Congruent neuronal modulation across competing actions challenges the role of the substantia nigra in action selection

Lingfeng Hou[1], Sofia Grijalva Torres[1,2], Daniel Svedberg[1], John Assad[1,*]

1 Department of Neurobiology, Harvard Medical School, Boston MA 02115

2 Department of Biological Sciences, the University of Texas at El Paso, El Paso, TX 79968

* Corresponding author: john_assad@hms.harvard.edu

Abstract

The basal ganglia are involved in the control of movement, but their exact role is unclear. Paradoxically, most of the inhibitory projection neurons in the main output nuclei increase firing around the time of movement; only a small fraction decrease firing. This antagonistic activity pattern could subserve action selection, with the small "decrease" population selectively disinhibiting the desired movement, and the larger "increase" population inhibiting competing movements. The action-selection hypothesis makes an implicit assumption: neurons that decrease firing to disinhibit a specific action should increase firing to inhibit that action when a different action is desired. To test this hypothesis, we recorded projection neurons in the substantia nigra pars reticulata (SNr) of mice trained to alternate between two different types of movements. Many SNr neurons showed a "ramping" pattern of pre-movement firing-rate modulation, with most neurons increasing firing, consistent with previous findings. However, contrary to the action-selection model, the overwhelming majority of SNr neurons exhibited congruent modulation between the competing actions, either increasing or decreasing their firing

rates for both actions; only a small fraction of neurons exhibited opposite signs of modulation. We could not ascribe the congruent modulation to potential uninstructed movements. Our results are not easily reconciled with simple antagonistic mechanisms for action selection in the basal ganglia output nuclei. We also found that ramping activity in SNr neurons typically began hundreds of ms before self-timed and spontaneous movements, in contrast to previous findings suggesting that basal ganglia output is modulated too late to be involved in movement initiation. Our findings suggest constraints – and raise new questions – about the role of the basal ganglia in movement initiation and action selection.

Introduction

The basal ganglia (BG) are a set of evolutionarily conserved sub-cortical nuclei involved in movement, motivation and other cognitive functions. The major input nucleus of BG, the striatum, receives extensive cortical and thalamic input, and in turn projects (directly and indirectly) to the major output nuclei of BG, the globus pallidus interna (GPi) and substantia nigra pars reticulata (SNr). The GPi and SNr contain tonically active GABA-ergic projection neurons which innervate thalamic and brainstem targets (Deniau et al., 2007; Nambu, 2007). Ascending projections to thalamic motor nuclei and descending projections to brainstem premotor nuclei in principle allow the basal ganglia to exert control over movement (Arber & Costa, 2022; Frost-Nylén et al., 2024; McElvain et al., 2021), but the mechanisms by which basal ganglia output influences movement remain unclear.

In the oculomotor system, the tonic discharge of SNr projection neurons provides continuous inhibition to the superior colliculus (SC) between saccades (Hikosaka & Wurtz, 1983a). At the initiation of saccades, many SNr neurons pause their tonic firing (Hikosaka & Wurtz, 1983b; Joseph & Boussaoud, 1985); this release of inhibition could allow a burst of collicular activity that triggers eye movement. In this view, basal ganglia output could gate movement via *disinhibition* (Deniau & Chevalier, 1992; Hikosaka & Wurtz, 1983c). However, a longstanding and ubiquitous observation is that the majority of SNr and GPi projection neurons *increase* their firing rates with movement. Neurons with movement-related increases in firing have been reported in rodent SNr (Bryden et al., 2011; Fan et al., 2012; Gulley et al., 2002; Lintz & Felsen, 2016; Rossi et al., 2016), cat SNr (Joseph & Boussaoud, 1985; Joseph et al., 1985), monkey SNr (Handel &

Glimcher, 1999; Magariños-Ascone et al., 1992; Sato & Hikosaka, 2002; Schultz, 1986; Wichmann & Kliem, 2004) and monkey GPi (Aldridge et al., 1980; Anderson & Horak, 1985; Brotchie et al., 1991; Georgopoulos et al., 1983; Mink & Thach, 1991a; Mitchell et al., 1987; Mushiake & Strick, 1995; Turner & Anderson, 1997). In most reports, movement-related “increase” neurons outnumber “decrease” neurons in the SNr and GPi, often vastly so (e.g., Bryden et al., 2011). These projection neurons are GABAergic – why would the basal ganglia increase their inhibitory output at a time when neurons throughout the motor pathways become more active? One possibility is that, simply by adding a sign change, “increase” neurons could activate downstream targets in parallel with decrease neurons. For example, increase neurons could preferentially innervate inhibitory neurons in target structures (Kaneda et al., 2008), or could secondarily excite downstream neurons through post-inhibitory rebound depolarization (Villalobos & Basso, 2022). In this view, increase and decrease neurons would act congruently to disinhibit their synaptic targets.

Alternatively, a longstanding hypothesis is that increase and decrease neurons act *antagonistically,* to focus or *select* actions (Alexander & Crutcher, 1990; Mink, 1996; Mink & Thach, 1993). For example, Mink and Thach (1991a) reported that 71% of movement-modulated neurons in monkey GPi increased their firing rate around the time of a visually cued wrist flexion or extension, whereas only 29% decreased. In addition, the distribution of the latencies of neural modulation (increase or decrease) largely overlapped the latencies of electromyographic (EMG) activity in agonist muscles. Given the preponderance of “increase” neurons – which presumably inhibit downstream targets – and the relatively late timing of neural modulation, the authors suggested that

GPi output was unlikely to play a role in movement initiation. Rather, the decrease neurons might select the *desired* movement by disinhibition, whereas the increase neurons might inhibit *competing* movement plans (Mink, 1996; Mink & Thach, 1991a). This could provide a sharpening of movement selection, analogous to how antagonistic center-surround receptive fields produce spatially sharper responses in sensory neurons (Alexander & Crutcher, 1990; Mink & Thach, 1993). Moreover, for any desired movement, one could imagine a far larger number of potential competing movements; thus, more SNr/GPi neurons should increase their activity than decrease – as is commonly observed.

Subsequent findings in the basal ganglia are often interpreted in light of the action-selection/focusing hypothesis. For example, co-activation of direct and indirect pathway striatal spiny projection neurons (dSPNs and iSPNs) has been suggested to play a role in action selection, whereby pro-movement dSPNs could promote the desired movement and anti-movement iSPNs could suppress unwanted movements, presumably through their modulation of SNr/GPi (Chen et al., 2021; Cruz et al., 2022; Cui et al., 2013; Gerfen, 2023; Klaus et al., 2019; Markowitz et al., 2018).

The action-selection hypothesis is intriguing, but lacks direct supporting evidence. In particular, “increase” and “decrease” neurons have only been defined functionally; nothing is known about whether they differ in terms of inputs, outputs, genetic identity, etc. However, a parsimonious view of the model is that increase and decrease neurons are two sides of the same coin: an SNr/GPi neuron that *decreases* its firing to specifically disinhibit movement *X* when it is desired could *increase* its firing to suppress

movement *X* when different actions are desired (Schroll & Hamker, 2013). This *bidirectional* response pattern could be viewed as analogous to that of sensory neurons, which are activated or suppressed depending on whether stimuli fall within the center or surround of their receptive fields. To address this question, a main goal of this study is to examine the activity of SNr neurons during different types of movements: do different movements elicit responses of opposite sign in SNr neurons?

The timing of movement-related GPI/SNr activity also bears closer scrutiny. During arm movements in monkeys, the majority of GPi neurons change their firing rate *after* the earliest arm-muscle EMG activity (Anderson & Horak, 1985; Mink & Thach, 1991a). Moreover, selective lesion or reversible pharmacological inactivation of monkey GPi do not increase reaction times (Desmurget & Turner, 2008; Horak & Anderson, 1984; Mink & Thach, 1991b; Turner & Desmurget, 2010). For these reasons, it has been argued that the basal ganglia do not play a primary role in movement initiation. However, in those studies latencies were measured in animals moving as quickly as possible in response to a sensory cue. Reactive movements of this sort are typically less impacted in Parkinson's disease (PD); PD patients have more difficulty with movements that are self-initiated (Glickstein & Stein, 1991; Martin, 1967; Russo et al., 2022; Sarma et al., 2012). Nonetheless, few studies have examined self-initiated movements in experimental animals. To address these questions, our lab introduced a controlled *self-timed* movement task (Lee & Assad, 2003). In the task, a sensory cue initiates a criterion timing period rather than triggering a rapid reactive movement; the animal must wait and then make the required movement *on its own*, without any further external cue. Previous findings from our lab and other labs suggests that striatal SPNs, motor cortical

neurons and dopaminergic neurons of the substantia nigra pars compacta (SNc) are indeed active early enough to play a role in *initiating* self-timed movements – exactly the types of movements most affected in PD (Hamilos et al., 2021; Lee & Assad, 2003; Lee et al., 2006). Thus, an ancillary goal of this study is to examine the activity of mouse SNr neurons during self-timed and self-initiated movements.

Results

Animal behavior in the self-timed reach task

Our main goal was to compare the activity of SNr neurons during two different types of self-timed movements, *reaching* vs. *licking*. We initially trained the first cohort of animals in the reach task alone. A second cohort was subsequently trained to perform both reach and lick tasks in interleaved blocks of trials. We shall first examine the self-timed reach task.

We trained water-restricted mice to perform the self-timed forepaw reach task while head-fixed (Figure 1a). At the start of each trial, a servomotor-mounted aluminum touch bar was swung into reaching range (800 ms travel time). With the touch bar deployed, a 100-ms start-cue (6,272 Hz tone) was played to indicate the start of the self-timing interval (Figure 1b). The task required the animal to wait at least 4 seconds after the cue and then touch the bar with its forepaw. No further cue was delivered after the initial cue; thus, the animals had to self-time the movement. If the touch occurred $\geqslant$4s after the cue (correct trial), a juice drop was delivered while the lick spout (normally retracted) was advanced via servomotor (800 ms travel time) into range of the tongue. The 800

ms delay in reward delivery was designed to discourage anticipatory licking while reaching. Extra touches made after the first did not yield additional reward. Both the touch bar and lick spout were retracted at trial end (10 s after cue). A touch before the 4s criterion time (incorrect trial) was not rewarded and resulted in immediate bar retraction followed by time-out until the end of trial (10 s after cue). If the animal did not touch the bar (no-move trial), it was retracted at trial end (10 s after cue). Thus, regardless of outcome, all trials ended 10s after the cue.

In the first training session, animals were rewarded for immediately touching the bar. The reward criterion-time was gradually raised from 0.1 s to 4 s over the first 3-5 sessions. Animals learned to delay their movements as training progressed (Figure 1c). We evaluated the behavioral performance by inspecting the distribution of movement times for each session. Most trained animals generated a Gaussian-like distribution with a mode ~4 seconds after the cue (Figure S1). Interval estimation in animals typically shows considerable variability around the criterion time (Gallistel & Gibbon, 2000; Hamilos et al., 2021; Lee & Assad, 2003), in contrast to reactive movements, which occur with a short, stereotypical latency after the “go” cue (Lee & Assad, 2003). Thus, we focused on trials where animals moved $\geqslant$2 s after the cue, since these are unlikely to be reactive movements prompted by the cue.

<u>SNr activity in the self-timed reach task</u>

We recorded single-unit extracellular activity in SNr using either chronically-implanted microwire bundles (adapted from (Starkweather et al., 2017)) or acutely-

implanted multi-electrode silicon probes (Yang et al., 2020 Masmanidis, 2020). Recordings were obtained unilaterally in 29 animals (17 right, 12 left) while they performed the self-timed reach task. To exclude dopaminergic neurons which are also sparsely located in SNr (typically discharging at 6 sp/s; Fan et al., 2012), we selected well-isolated units with a minimum baseline firing rate of 15 sp/s as putative GABAergic SNr neurons. 784 units were considered putative SNr projection neurons (median baseline spike rate 39.7 sp/s, median absolute deviation 11.7 sp/s). (The number of animals, recording sessions, and putative SNr neurons involved in the self-timed reach task, as well as subsequent experiments, are listed in Table S1.)

During the post-cue/pre-movement period leading up to the forepaw reach, many SNr units either increased or decreased their firing rate. The pre- and peri-movement response time-courses of two SNr units are shown in Figure 1d, aligned to the time of touch-bar contact. In both examples, and in the majority of SNr neurons that we recorded, the average pre-movement activity changed gradually, starting hundreds of milliseconds before the movement. This pre-movement pattern of activity was reminiscent of “ramping” activity that our lab previously observed for self-timed movement tasks, in striatum (Lee & Assad, 2003; Lee et al., 2006), parietal cortex (Maimon & Assad, 2006) and SNc (Hamilos et al., 2021).

Peri-event time histograms (PETHs) of all 784 SNr units are shown as a heatmap in Figure 1e. The “baseline” spike rate of each SNr unit was estimated over the [-4 s, -2 s] window before bar contact. Because the forepaw was already moving by 200 ms before the touch-bar contact (see Figure 2), we calculated “peri-movement” activity over the [-0.3 s, 0] window before bar contact, which generally captured the peak change in

activity for individual units. Using an early-shifted “pre-movement” window of [-0.5 s, -0.2 s] yielded similar results in all subsequent analyses.

We measured the relative amplitude of peri-movement responses by subtracting the baseline rate (Figure 1g), or by normalizing to the mean and standard deviation of the baseline rate (Figure 1h). 641/784 (82%) SNr units showed significant pre-movement spike-rate modulation relative to baseline (permutation test, $p<0.01$; see Methods). Of these 641 units, 496 units (77%) increased pre-movement firing and 145 (23%) decreased pre-movement firing. The prevalence of “increase” units over “decrease” units is similar to that described in many previous studies in SNr/GPi (e.g., Mink & Thach, 1991a).

The decrease population had a significantly higher median baseline spike rate (44.6±12.1 sp/s) than the increase population (39.7±11.0 sp/s; one-sided Wilcoxon rank-sum test; $p<0.0023$). Nonetheless, there was a broad overlap in the distributions of baseline-firing rates (Figure 1f), and baseline rates were not a reliable basis for classifying increase vs. decrease SNr units (e.g., area under the receiver-operator curve = 0.58).

SNr activity modulation precedes movement onset

Many SNr units showed spike-rate modulation starting hundreds of milliseconds before movement, manifesting as a gradual downwards (or upwards) “ramp” in the activity of individual SNr units, as well as in population-average activity (Figure 2a). For the 641 significantly modulated units, we quantified the time of onset of neural

modulation relative to the time of touch-bar contact by determining when trial-averaged spike rates first deviated ⩾0.25 SD from the mean continuously for 50 ms (see Methods). The distribution of response-onset times was broad, with many units showing modulation arising hundreds of ms before the touch-bar contact (Figure 2b). 98% of SNr units had a response onset >200 ms before touch-bar contact, and the tail of the distribution extended beyond 2000 ms. The median response-onset time was -1104 ms for increase units, and -615 ms for decrease units. However, these times were relative to touch-bar contact, but animals had to first lift their forepaws and reach before contacting the touch-bar. To determine the time of forepaw-movement initiation, we used markerless pose estimation from video data (DeepLabCut; see Methods). Among 1,414 self-timed reach trials (11 animals, 22 recording sessions), the median movement-onset time was 175 ms before touch-bar contact, and 90% of forepaw movements were initiated within 450 ms before contact (Figure 2c). The median neural-onset times were significantly earlier than the median movement onset time, for both increase and decrease populations ($p < 8.7\times10^{-194}$ and $p < 2.4\times10^{-54}$, respectively; one-sided Wilcoxon rank-sum test). We also measured neck-muscle EMG activity in one animal (3 sessions, 336 trials). The forepaw reach reliably evoked a signal in the neck muscles. On average (across 336 trials), EMG activity onset (when EMG activity exceeded 0.25 SD for 2 consecutive 25-ms sample points; see Methods) preceded video-detected movement-onset (when forepaw displacement exceeded 0.25 SD for 2 consecutive 25-ms sample points; see Methods) by 199 ms. Even when we artificially lengthened the (video-detected) forepaw movement onset latency by 200 ms to account for the EMG offset, the median neural-onset times still significantly preceded the median

movement-onset time ($p < 1.2\times10^{-136}$ and $p < 2.1\times10^{-18}$ for increase and decrease units, respectively; one-sided Wilcoxon rank-sum test), and many neurons showed firing-rate modulations that preceded the earliest EMG activity by hundreds of milliseconds.

The early onset of firing-rate modulation relative to movement suggests that SNr neurons could play a role in initiating self-timed movements. However, SNr and other motor neurons respond to “jittery” spontaneous body movements on a much shorter timescale (Hasnain et al., 2025; Rossi et al., 2016). We indeed observed phasic SNr responses to spontaneous limb, torso, and orofacial movements during intertrial interval. Although our subjective impression was that mice tended to “settle down” at the start of self-timed movement trials, it was possible that the animals made more frequent jittery movements as trials progressed, resulting in a seemingly gradual ramping of neural activity in SNr.

To examine the possibility that SNr activity could be explained by jittery movements alone, we collected high-quality bilateral video data for 6 animals (7 sessions, 163 SNr units recorded), and used markerless video tracking to determine the position and velocity of all four paws, the nose, and the midpoint of the spine (see Methods). When averaged across trials, the velocity of all body parts was close to zero until the initiation of the forepaw reach (Figure S2c). This suggested that the ramping neural modulation which preceded movement initiation was likely not due to a gradual acceleration of jittery body movements.

To examine this question on a trial-by-trial basis, we constructed a series of nested generalized linear models, i.e., Poisson regression (Pillow, 2007), using four categories of predictors to estimate spiking rates for individual SNr neurons. 163 SNr units were

modelled independently on a unit-by-unit basis. “Cue” (C) predictors were discrete “trial-start” events convolved with 400-ms wide cosine kernels with temporal offsets ranging from -600 to 600 ms (negative offsets were included because touch-bar deployment, which was audible, began 800ms before the cue). “Velocity” (V) predictors included video-derived movement velocities of the four paws, the nose, and the midpoint of the spine, with temporal offsets ranging from -200 to 200 ms. “Ramp” (R) predictors included a series of linear ramping signals with different slopes, all terminating at the time of touch-bar contact but with different onset times ranging from 100ms to 2000ms before touch-bar contact. Finally, the “trial-progression” (T) predictor was a single linearly increasing signal beginning at cue onset and ending at touch bar contact.

Using a nested approach, we iteratively added predictors to the GLM and asked whether they improved model performance (Figure 2d), as measured by variance explained (R^2) or Akaike information criterion (AIC). The “null” model was a constant reflecting the average spiking rate across the whole session. The “null + cue” (+C) model accounted for cue-responses in a small number of units, but performed poorly overall. The “null + cue + velocity” (+V) model could predict the sign and amplitude of pre-movement modulation, but failed to predict the early onset of the ramp (Figure 2f). In contrast, the “null + cue + velocity + ramp” (+R) model not only explained additional variance (Figure 2d), but also allowed for significantly better reconstructions of the pre-movement PETH (Figure 2e shows reconstructions from two example SNr units). Finally, the addition of the “trial-progress” (+T) signal did not further improve model performance (Figure 2d). Thus, we did not find evidence that a gradual ramp-up of jittery random movements could account for the early onset of spike-rate modulation in

SNr neurons. Rather, the start of the ramping activity was related to the initiation of the upcoming self-timed movement in some way.

Slope of ramping activity in SNr is timing-independent

Previous work from our lab showed upward-ramping activity among dopaminergic neurons in the SNc within a self-timed lick task. The average ramping activity began at the start-timing cue and continued until the time of movement. The slope of the ramping varied according to (and was predictive of) when animals chose to move on a trial-by-trial basis, whereby steep slopes were associated with early movements, and shallow slopes with later movements, as if the neural activity rose toward some systems-level threshold (Hamilos et al., 2021).

In contrast to our previous findings in the SNc, the onset and slope of ramping activity in SNr neurons was much more stereotyped when aligned to the start of movement (Figure 2g). To quantify this effect for each SNr unit, we grouped trials into 4 bins based on movement time (2-4, 4-6, 6-8, and 8-10 seconds) and then constructed movement-aligned PETHs for each of the four movement-time bins. For each neuron, we averaged the Euclidean distances between each pair of PETH traces (six pairwise comparisons for four bins) during the pre-movement period ([-2, 0] s prior to bar contact (to avoid reward-related differences that arose following the movement). We compared that average distance to the permuted distribution of average distances obtained from repeatedly shuffling trial identity (see Methods). Only 30/641 (5%) units showed significant differences among the timing-dependent PETHs ($p < 0.01$). Even for those few units, individual inspection revealed that detectable differences were most

prominent at the beginning of trials, likely reflecting cue-related responses on short trials rather than pre-movement activity (Figure S2d). In conclusion, we did not find evidence that ramping activity of SNr neurons varied with the timing of the movement. Rather, the ramping activity for each SNr unit generally exhibited a stereotyped slope and onset when aligned to the start of movement.

SNr activity precede spontaneous movements

For both increase and decrease units, we found robust ramping activity (on average) arose hundreds of ms before the earliest detection of forepaw movement. This early neural modulation suggests that the SNr could, in principle, be involved in movement initiation. However, ramping activity could be an idiosyncratic feature of the timing requirement of our self-timed movement task and might not be present for "spontaneous" movements. In particular, the self-timed movement task required animals to *withhold* movement until the criterion time. Firing of inhibitory SNr neurons might play a role in active withholding, similar to how striatal iSPNs activated during a "dynamic action suppression" task (Cruz et al., 2022). At least for the increase neurons, the gradual ramping up in activity might counteract the animal's increasing urge to move throughout the timing interval.

Several observations argue against this possibility, First, if ramping activity of increase units acts to withhold movement, we might expect that their activity should abruptly *decrease* just before movement initiation, as if withdrawing the brake. However, the SNr units in our sample generally showed monophasic responses: the sign of the peri-movement response was consistent with the sign of the ramping activity (Figure 1e;

Figure S2a). Second, at the start of each trial in the reach task, the animals consistently made an overt “flinching” movement as the servomotor began to deploy the touch-bar. This flinching often evoked a phasic response from SNr units. Invariably, the sign of this phasic cue-response matched that of the ramping activity (Figure S2a & b): units that ramped up/down during the self-timed interval generally showed a phasic increase/decrease in activity in response to the overt flinch, respectively. These combined observations suggest that the ramping activity during the self-timed movement task was related to the upcoming movement *per se* rather than *withholding* that movement.

Notwithstanding, it is still possible that the ramping activity is a peculiar feature of self-timed movements and is not a general feature of “spontaneous” movements. To examine this question, we designed a spontaneous reach task that eliminated the explicit timing requirement (Figure 3a). Animals were allowed to move at any time, and touching the bar always resulted in a rapid “cycling” of the touch bar – immediate retraction followed by immediate redeployment – requiring a total of 1.4 s. A randomized timeout period started after any rewarded touch. During the timeout, reaches were not rewarded and resulted in immediate cycling of the touch-bar. The end of the timeout interval was not signaled to the animal. The first reach after the timeout was rewarded. The lick-tube was presented for 3-4 s for reward collection, followed by the next timeout interval. Importantly, the duration of the timeout period was randomly drawn from an exponential distribution, with a time constant of 20 s. The exponential probability was chosen because the instantaneous probability of the event of interest does not change with elapsed time (i.e., flat hazard function).

We trained three naïve mice to perform the spontaneous reach task. Across all 8 sessions in the 3 animals, the distribution of inter-reach intervals declined monotonically, resembling an exponential distribution (Figure 3b) rather than the peaked distribution we observed for animals trained in the self-timed movement task. This is consistent with the animals not using a timing strategy.

We recorded from 89 putative SNr neurons in the 3 animals. The pre-movement modulation of spiking activity is shown in Figure 3c-e. A subset of 47/89 units exhibited significant pre-movement modulation relative to baseline ($p<0.01$, permutation test): 41 units (87%) increased in firing rate, and 6 units (13%) decreased in firing rate (Figure 3d-e). On average, the pre-movement modulation of SNr activity for spontaneous reaches resembled a gradual ramp, with most units increasing activity and a minority decreasing activity, very similar to the ramping activity we observed in the self-timed movement task (compare with Figure 1). This suggests that the pre-movement ramping is not specific to explicit timing tasks, and thus that ramping activity is unlikely to represent timing-related active withholding or suppression of movement.

Interleaved reach vs. lick task to test the sign of neural responses in SNr

In the self-timed reach task, the majority of modulated units increased their firing rate during the pre-movement period, while only a minority of units decreased their firing. According to the movement-selection hypothesis, the small contingent of decrease neurons could be selective for the forepaw movement, with their decreased firing disinhibiting the reaching movement. It follows that the larger population of increase neurons (during reaching) were selective for a different type of movement not

involved in forepaw reaching. Their increased firing would thus have inhibited those unwanted movements during reaching. In the simplest view of this action-selection scenario (Mink & Thach, 1993), if animals were to perform an alternative movement – licking, for example – the neurons that previously decreased firing during reaching would now increase their firing, so as to now inhibit forepaw movements, while a subset of the neurons that previously increased their firing during reaching – i.e., those that are selective for licking – would now decrease their firing to disinhibit licking. In this view, some individual SNr neurons should show *opposite signs* of response during the two types of movements (Schroll & Hamker, 2013).

To test this hypothesis, we trained a subset of 16 animals to perform two types of self-timed movements – forepaw reach vs. lick – in interleaved blocks of 50-75 trials within the same session. The reach task was as described above (Figure 1b), with the lick-spout remaining retracted until after a correct reach was made, to discourage anticipatory/spontaneous licking. The lick task (Figure 4a) largely resembled the reach task except for two key differences: 1) the touch-bar was retracted out of reach throughout the trial, and the lick spout was deployed (instead of the touch bar) at the start of trial; and 2) animals were rewarded if the first *lick* occurred 4-10 seconds after the start cue. In both tasks, servo motors deployed or retracted the touch-bar or lick-spout over 0.8 s. For both tasks, animals were only rewarded for the first correctly timed movement (reach or lick, depending on the task) in each trial; moving more than once did not yield more rewards. Early contact (before 4 s) with the touch bar or lick spout resulted in their immediate retraction and a timeout period. All trials (regardless of

outcome) were 10 seconds in length, so there was no advantage to making early movements.

Naïve animals were initially trained to perform only the reach (or lick) task, then learned to perform both reach and lick tasks in interleaved blocks of 50-75 trials within the same session. Within the same session, well-trained animals were required to perform at least 20 self-timed reach trials and 20 self-timed lick trials for which the movement occurred $\geqslant$2 s after the cue. Animals reached asymptotic behavior within 7-14 days of training (Figure 4b), exhibiting a broad distribution of movement times in both the reach and lick tasks (Figure 4c-d; Figure S4).

Importantly, we designed the interleaved reach and. lick tasks to elicit an “isolated” reach or lick movement, by 1) by retracting the touch-bar out of reach throughout lick trials; and 2) delaying lick-spout deployment on reach trials. Nonetheless, animals could have made anticipatory licks before or while they reached (even though the lick tube was retracted during the reach) or fidgeted their forepaw before they licked. If the animals made *both* movements on each trial – one instructed and one spurious – this would undermine the goal of the experiment. To control for this, we collected and analyzed high resolution bilateral video recordings for a subset of 6 animals (7 total sessions). Trial-averaged time courses for movement speeds of the contralateral forepaw, the spine, as well as probability of licking/tongue protrusion (including “air licks” that failed to reach the spout), are shown in Figure 4e. During the reach task, animals generally did not engage in anticipatory licking before or during the forepaw movements: licking only occurred when the lick spout was moved into range of the tongue on correct trials (Figure 4e, top right). This delay in licking would presumably

allow us to distinguish neural activity related to the earlier reach from the later lick. Conversely, animals did not make detectable forepaw movements (on average) during the lick task (Figure 4e, bottom row). In addition, video-tracking did not reveal evidence of overt bodily movements (e.g., along the spine) that might be related to postural adjustments common to both movements (Figure 4e, green trace). These behavioral measurements suggests that we were able to successfully isolate licking vs. reaching between the two tasks, which should allow us in turn to distinguish reach-related from lick-related neural activity.

Activity during reach vs. lick task is largely congruent in individual SNr units

We recorded the activity of 427 SNr units from 16 well-trained animals. Figure 5a shows spike rasters from three example units with different response patterns: 1) decreased firing in both reach and lick trials; 2) increased firing in both reach and lick trials; 3) oppositely modulated firing in reach vs. lick trials. Across all 427 units, in both tasks, we observed more increase units than decrease units (Figure 5b). The majority of SNr units exhibited *congruent* modulation during reach and lick, either increasing or decreasing their peri-movement firing rate for *both* tasks. However, a small subset of the units showed opposite signs of modulation between the two tasks (Figure 5c).

To quantify these findings, we averaged the peri-movement response of each SNr unit over the window [-300, 0] ms before target contact, for both tasks, and we determined if the response was statistically significantly different from the baseline activity, as described above (permutation test, $p<0.01$). For the two tasks separately, we classified SNr units as “increase”, “decrease”, or “unmodulated”, based on their sign of

(or lack of) response. Figure 5d shows the average peri-reach vs. peri-lick responses for all 427 units. 272/427 units were significantly modulated in both tasks, allowing us to classify the sign of the response for both reach and lick (the remaining 155 units were significantly modulated in only one or neither of the two tasks). Of these, the vast majority (234/272, 86%) had congruent signs of neural modulation for both forepaw and licking movements, either increasing or decreasing firing for both tasks. Only a small minority of SNr units exhibited a sign reversal across the two tasks – 14% of units significantly modulated in both tasks; 9% of all recorded units. Across the entire population of 427 SNr units, there was a significant positive correlation between peri-reach and peri-lick modulation among all 427 SNr units (R^2=0.29, $p < 4.3\times10^{-34}$ under null hypothesis of zero-slope; linear regression), consistent with the majority of units having a congruent sign of modulation between the two tasks.

Of special interest are the subset of SNr units that showed decreased firing during the peri-movement period. Under the simple action-selection hypothesis, the decreased firing in these units could act to disinhibit a desired movement – but then those same units should increase firing to suppress that movement when it is becomes undesired. However, even among the 93 units that showed significant *decreases* in firing in either task, only a minority showed increased firing in the other task (38/93, 41%). In conclusion, few SNr neurons showed opposite signs of firing-rate modulation between the reach and lick tasks. This finding is difficult to square with a simple model whereby SNr neurons select or de-select actions.

Congruent SNr activity in the reach vs. lick task is not due to movement planning

Whereas most SNr units showed congruent signs of activity in the reach vs. lick task, with either increased or decreased firing in both tasks, we were initially concerned that spurious forepaw movements (during the lick task) or anticipatory licking (during the reach task) could artifactually present as congruent neural modulation. Using video-tracking, we detected neither spurious forepaw movements during the lick task nor anticipatory licking during the reach task (Fig. 4e), allowing us – at least in principle – to distinguish neural activity related to reaching and licking. However, even if *overt* movements were dissociated in time, animals could in principle still *plan* to perform both actions, and this could potentially explain the apparent congruent neural responses we observed between the two tasks. This was less of a concern for the lick task, because the touch bar was retracted out of reach of the animals throughout the entire block of lick trials: presumably, the animals did not plan to reach if there was nothing to reach to. The interleaved reach task, however, could be more confounded, because the animals could anticipate licking following the reach on correct trials. In this scenario, if a unit were selective only for licking, that unit might then show congruent pre-movement modulation between the lick and reach task, because the animal planned to lick in both cases; this could lead us to mis-characterize that unit as non-selective between the tasks. To mitigate this concern, we delayed the licking at the end of reach trials by slowly deploying the lick-tube (over 800 ms) into range of the tongue following a correct reach. We reasoned that by delaying the lick-tube deployment, potential lick-related pre-movement modulation would likewise be delayed for lick-selective units. Indeed, video tracking revealed that tongue movements were delayed in the reach task compared to

the lick task (Figure 4e, top vs. bottom row). However, the pre-movement neural response was *not* systematically delayed in the reach task; if anything, the onset of pre-movement neural modulation was noticeably *earlier* in the reach task than the lick task (Figure 5b-c). This suggests that the pre-movement neural modulation during the reach task was indeed related to the forearm reach, not planning the upcoming lick.

A second observation also argued for a distinct temporal separation between reach-related and lick-related neural activity on reach trials. On correct reach trials, animals typically started licking (to collect the juice reward) ~0.8 s after contacting the touch-bar, when the lick-spout was deployed within reach of the tongue. For the minority of neurons that had *opposite* signs of modulation for reaching vs. licking, when the animals began to lick on reach trials, those SNr units abruptly changed their firing rate to be consistent with the firing-rate modulation observed during the self-timed lick task. That is, units that had increased firing during the reaching phase rapidly decreased their firing below baseline level when licking commenced, and units that had decreased firing during the reaching phase rapidly increased their firing above baseline level when licking commenced (Figure 5e). Even for units with congruent responses in the two tasks, there was a distinct inflection in the neural response when licking began following correct reaches (Figure 5e). These findings suggest that the lick neural response was indeed delayed relative to the reach neural response. Following this logic, the congruent neural responses that we observed in most SNr neurons were likely unrelated to the mice “planning ahead” to lick on reach trials.

Amplitude of neural modulation in the lick vs. reach task

While the sign of neural modulation was congruent between lick and reach for most SNr units, we also examined the relative amplitude of peri-reach vs. peri-lick modulation. 188/427 (44%) units exhibited significantly different modulation amplitudes between reach and lick tasks ($p<0.01$; permutation test, see Methods), even though the signs of the responses were largely congruent between the two tasks. At the population level, these differences were substantive enough that we could use SNr responses to classify individual trials as reach or lick trials. We trained a linear-discriminant decoding model (see Methods) to predict the underlying trial type using population neural activity ($\geqslant$10 SNr units/session, 13 sessions, 293 total units; 490 reach trials vs. 680 lick trials). Although the models were trained on neural activity from the peri-movement period ([-0.3, 0] s prior to touch-bar or lick-spout contact), they were able to successfully decode the true movement type hundreds of ms before and after the movement (Figure 5f). Therefore, although individual SNr neurons mainly exhibited congruent signs of modulation for reach vs. lick, the population of SNr neurons still carried information about the type of movement on a trial-by-trial basis.

Lick-entrained activity modulation is common among SNr neurons

On rewarded trials, the lick spout was presented for 2-6 seconds (depending on movement time). During this period, animals performed bouts of rhythmic licking with a stereotypical frequency of 8-10 Hz (Figure 6a). We identified a sub-population of 261/427 (61%) SNr neurons whose firing was strongly entrained to the periodic protrusion/retraction of the tongue observed during lick bouts (Figure 6b; $p < 0.01$,

permutation test on circular statistics of spiking rates; see Methods). The phase of the entrainment varied considerably among the 261 units, covering the full period of the lick cycle. We therefore grouped the units into four categories according to phase: in-phase (increased spiking on spout contact), anti-phase (decreased spiking on spout contact), and first-quarter/third-quarter phases. The distribution of phase preference is shown in Figure 6c, along with their lick-triggered population-average spiking activity in Figure 6d.

The periodic firing pattern of these 261 lick-entrained neurons is presumably related to licking *per se*, rather than, for example, some gross postural adjustment. However, even for this subset of lick-entrained units, the signs of the peri-movement modulation between the lick and reach tasks were overwhelmingly congruent. We found a positive correlation in peri-movement modulation between the reach and lick tasks (R^2=0.33, $p < 4.7\times10^{-24}$; linear regression; Figure 6e). We further hypothesized that in-phase lick-entrained SNr units were more likely to increase firing during the ramp-up period before the first self-timed lick (and vice versa for anti-phase units). However, phase-preference during lick bouts did not predict peri-movement activity for either self-timed licks or self-timed reaches (Figure 6f). Rather, lick-entrained activity resembled small fluctuations in spike rates superimposed on the slower, higher-amplitude modulation observed during self-timed movements (Figure 6g). It is therefore possible that lick-entrained oscillatory activity typically associated with SNr (Rossi et al., 2016) is orthogonal to the envelope of self-timed movement activity in the lick and reach tasks. Regardless, the observation that even lick-entrained SNr units also showed largely congruent neural responses for reach *and* lick argues against the simplest action-selection model (see Discussion).

SNr responses do not change sign for different reach directions

We observed few SNr neurons with opposite modulation for reach vs. lick; however, it is possible that a putative antagonistic action-selection mechanism in the SNr would only apply to movements that are *mutually exclusive*. Arm and orofacial movements are presumably not mutually exclusive (in the sense that “walking and chewing gum” are not mutually exclusive). But *different* movements of the *same* effector are, by definition, mutually exclusive -- they cannot be performed at the same time. We thus designed a bi-directional movement task in which animals reached towards two different targets on interleaved trials. Animals were trained to use their contralateral forepaw to touch a small metal cube (5 mm wide by 10 mm high). Upon trial start, the target was presented at one of two possible locations, either lateral or medial to the contralateral forepaw. The target position was adjusted using a linear actuator before being deployed into range using a servomotor. Animals used their contralateral forepaw for most reaches; trials with reaches of the ipsilateral paw (or both paws) were excluded from further analysis.

Animals were trained to make self-initiated reaches in the spontaneous movement task described in Figure 3. The target position was changed after every 15 rewarded trials. We recorded the forepaw movements with three cameras (front/left/right) and manually traced and reconstructed the movement trajectories of both forepaws. The direction of reach trajectories of the contralateral forepaw indeed differed depending on target position, while the ipsilateral paw remained stationary (Figure 7a).

We recorded 179 SNr units (10 sessions) in 3 animals trained to perform the bi-directional reach task. The onset of arm movement was determined through video

tracking. Reaches that followed a previous reach by <4s were excluded to focus on self-initiated rather than repetitive movements. For each SNr unit, we calculated the normalized trial-averaged peri-movement activity time course during lateral reaches vs. medial reaches (Figure 7b). For visualization and comparison, spike rates were normalized to the mean and SD during the [-4 s, -2 s] baseline window prior to reach onset.

Similar to the single-target reach tasks described above, we observed early neural modulation, as well as a majority of units that increased rather than decreased their pre-movement activity (Figure 7b). However, peri-movement neural modulation in SNr (within the window [-100 ms, 200 ms] relative to reach onset) was largely similar between the two reaching targets. 95/179 (53%) SNr units exhibited significant ($p<0.01$; permutation test) peri-movement modulation for one or both reach directions, but none of them (0/95) exhibited opposite *signs* of modulation (Figure 7c). Furthermore, only 10/179 (6%) exhibited a significant difference in modulation *amplitude* ($p<0.01$; permutation test) between lateral vs. medial reaches.

We then asked whether activity in SNr could be used to decode lateral vs. medial reaches made with the contralateral forepaw. We performed linear discriminant analysis by training decoding models to predict whether the animal was making a lateral or medial reach, using population neural activity (7 sessions with $\geqslant$10 units/session, 165 total SNr units; 351 lateral vs. 424 medial reach trials). The trained models were not able to distinguish lateral vs. medial reaches: decoding performance was not significantly better than chance ($p>0.01$; comparing to null models trained after shuffling the trial labels; 10,000 permutations; Figure 7d). Similar results were found for a 4-

direction reach task, with 3 contralateral targets and 1 ipsilateral target, with one noticeable difference: SNr neurons exhibited slightly higher firing rates (and thus presumably produced more inhibitory output,) when reaching with the ipsilateral forepaw, presumably a manifestation of the contralateral bias of forepaw motor control (Figure S7).

In conclusion, SNr neurons did not change sign or amplitude of modulation for reaches made with the same forepaw towards different directions, despite the mutual exclusivity of those actions.

SNr functional somatotopy is not clearly defined

We found few neurons with opposite response signs for different movement types in SNr, but there may be some order in the response types if we consider the anatomical location of individual neurons within the SNr. The SNr receives convergent input from striatum: anatomical viral tracing studies (Foster et al., 2021; Hintiryan et al., 2016) have revealed seemingly parallel, spatially segregated pathways within striatonigral projections, with the (dorsolateral) limb region of striatum projecting to ventromedial SNr, and the (ventrolateral) orofacial region of striatum projecting to dorsointermediate SNr (Foster et al., 2021; Lee et al., 2020).

We used 4-shank silicon microprobes (Yang et al., 2020 Masmanidis, 2020) to record throughout the depth of the SNr, in a spatially selective manner. The probe contacts spanned an area 0.45 mm along the medial-lateral axis and 0.78 mm along the dorsal-ventral axis. The ventral tip of the probes targeted the presumed limb region of

SNr, while the dorsal half of the two lateral shanks targeted the putative orofacial region (Figure 8a). In the self-timed reach vs. lick task, a subset of 242 putative SNr neurons (8 animals, 12 sessions) were collected using these probes (Figure 8b). As described above, we performed permutation tests to identify units with significantly positive or negative signs of peri-movement modulation ($p<0.01$). We also identified units with significantly different amplitudes of modulation for reach vs. lick ($p<0.01$).

Assuming that neurons that decrease their firing rate for a given movement type will disinhibit that movement, we would expect that *reach-decrease* and *lick-decrease* units would congregate in the ventromedial and dorsolateral portions of our recording sites, respectively. We visualized the approximate stereotaxic coordinates of 193 units significantly modulated during the reach task (Figure 8c) and 186 units modulated during the lick task (Figure 8d). Contrary to our prediction, we found no evidence of segregation of reach- or lick-modulated units in the SNr. Furthermore, 23 SNr units with opposite signs of modulation for lick vs. reach (18 reach-decrease/lick increase and 5 reach-increase/lick-decrease; Figure 8e & f), as well as 81 units that exhibited significantly different response *amplitudes* for lick vs. reach (Figure 8g), were also distributed across the four shanks.

Lick-entrained neurons had been previously observed in the dorsointermediate portion of SNr (Rossi et al., 2016), which preferentially receive input from ventrolateral (i.e. orofacial) striatum (Foster et al., 2021; Lee et al., 2020). However, 161 lick-entrained SNr units were evenly distributed throughout all four shanks of our recording probes (Figure 8h), including but not limited to the dorsointermediate part of SNr. Thus we did not find evidence for clear *functional* somatotopy within the SNr, in contrast to

the segregated striatonigral circuits that have been defined by monosynaptic viral tracing.

Discussion

We made three main findings. First, many SNr neurons modulated their firing rate before self-timed movements, with most neurons increasing rather than decreasing firing. Second, between different types of movement, the vast majority of SNr neurons had congruent signs of modulation, either increasing or decreasing; only a small minority of SNr neurons showed opposite signs of modulation. Third, for many neurons the modulation in firing rate began at least hundreds of milliseconds before the onset of self-timed and spontaneous movements. These observations suggest constraints – and raise new questions – about the role of the basal ganglia in movement initiation and action selection.

Relating our findings to the action-selection model

Since the 1980s, researchers have observed a paradoxical pattern of activity in basal ganglia output neurons during behavior. First, many studies have reported that the majority of SNr and GPi neurons increase rather than decrease their firing rates with movement. We likewise found the preponderance of SNr neurons that increased their activity leading up to movement, in multiple experimental contexts that we tested. SNr and GPi output neurons are GABAergic and tonically active; why would the basal ganglia relay even more inhibition to downstream targets around the start of movement? Second, most GPi neurons were previously reported to activate concurrently with or

after the onset of EMG activity, seemingly too late to play a role in movement initiation. Yet hypo- and hyper-kinetic movement disorders suggest a role for the basal ganglia in movement initiation. How can the basal ganglia contribute to movement initiation if their output modulation does not precede movement?

To reconcile these observations, more than 30 years ago Mink and Thach proposed that basal ganglia output neurons could be involved in *selecting* rather than initiating movements (Mink & Thach, 1991a; Mink & Thach, 1993). They hypothesized that the majority of pallidal output neurons increase firing to inhibit unwanted actions, whereas a minority decrease firing to disinhibit the specific action desired at that moment. They further speculated that structured input from the antagonistic striatal direct and indirect pathways could shape antagonistic neuronal responses in the GPi/SNr (Mink, 1996). Indeed, subsequent findings in the striatum, such as co-activation of striatal dSPNs and iSPNs during movement, have been interpreted in this light (reviewed in Bariselli et al., 2019).

We set out to test this model by recording from SNr projection neurons in animals trained in self-timed or spontaneous movement tasks. We chose these behaviors based on clinical observations that Parkinson's patients have more difficulty generating self-initiated movements compared to externally cued movements. We specifically tested an implicit prediction of the model, that SNr neurons should show a *reversal* in the *sign* of their modulation for different movements. That is, neurons that decrease firing to disinhibit a specific movement, *X,* should increase their firing to suppress movement *X* when the animals makes a different movement. However, among the hundreds of SNr neurons recorded in our experiments, only a small fraction showed this sign-reversal

phenomenon. This was the case whether the animals alternated between reach and lick movements (9% of neurons), or reached to medial vs. lateral targets (0% of neurons). Rather, the vast majority of SNr units evinced sign-congruent responses for the different types of movements, either increasing or decreasing firing for *both* movements.

Potential explanations for the paucity of sign-inverting SNr neurons

At first glance, our findings seem largely incompatible with a simple action-selection model in which SNr neurons select/suppress actions by inverting the sign of their peri-movement responses. However, there are several hypothetical scenarios whereby we could have missed sign-inverting neuronal responses or over-estimated the prevalence of sign-congruent neuronal responses. Consider first the neurons that decreased firing for both types of movements. It is possible that both tasks engaged a common movement that was not explicitly instructed yet still necessary for both tasks, such as a postural adjustment of the body. According to the action-selection model, if a neuron were selective for that common movement, it should decrease firing to disinhibit that movement in both tasks. Second, consider the SNr neurons that increased firing for both movements, e.g., reach and lick. These neurons might be selective for a third, unrelated movement that is not involved in either licking or forepaw reaching, such as movement of the hind-paw or tail. By increasing firing during licking *and* forepaw reaching, the congruent increase neurons would thus suppress that undesired movement during both tasks. Under these two scenarios, our finding of congruent (non-inverting) modulation could still be compatible with the action-selection hypothesis. In

this view, had we tested a broader range of movements, we might have found more neurons with sign reversals.

However, there are several arguments against these hypothetical scenarios. First, licking and reaching are distinct mechanically, and would thus likely involve different compensatory postural adjustments. We also video-tracked other parts of the body (e.g., spine) during licking and reaching and did not observe any systematic common movements. Thus, we did not uncover evidence for a “third”, common movement that might otherwise explain the responses of congruent decrease neurons.

Second, these scenarios would imply that the SNr neurons with congruent responses – the vast majority in our study – were selective for some (unknown) movements *other than* reaching or licking. However, representations of the contralateral forearm/paw and orofacial are expansive in somatotopic motor maps in rodents (Hintiryan et al., 2016), and we deliberately targeted electrode penetrations to sample from areas of the SNr reported to represent forepaw and orofacial movements (Foster et al., 2021; Lee et al., 2020). Thus, it is unlikely that we under-sampled neurons that were selective for forepaw reaching and licking, especially given the large sample size in our study.

Apropos, many neurons in our sample showed clear cycle-by-cycle entrainment in their pattern of firing during the repetitive lick bouts on rewarded trials. These neurons could be involved in protruding or retracting the tongue, or in modulating other periodic orofacial movements intrinsic to collecting and consuming the liquid (e.g., jaw movement, swallowing, etc.). By any definition, this pattern of neuronal activity is “lick-related”, yet the vast majority of these lick-related neurons still showed congruent

activity between licking and reaching, either increasing or decreasing their firing rate before both types of movement. The reach task involved a sequence of reaching first and then licking after a delay, yet we found no evidence that the pre-movement ramping activity on reach trials was due to the upcoming lick; on the contrary, activity related to licking was delayed commensurate with the ~1s delay between reaching and licking on those trials (Figure 5e). Under the action-selection hypothesis, we should not have observed lick-entrainment among the “congruent-increase” neurons, because they would have encoded some other movement(s) not involved in *either* reaching or licking; nonetheless, the vast majority of lick-entrained neurons showed congruent modulation between the lick and reach tasks (Figure 6e).

Finally, it is possible that an antagonistic action-selection mechanism in the SNr would not apply to reaching vs. licking, because those movements are not inherently mutually exclusive. To test this hypothesis, we trained animals in the directional reaching task: reaching laterally or medially with the same forepaw cannot happen at the same time. Even for this task, however, we found overwhelmingly congruent peri-movement modulation.

Given these contradictions, it is difficult to square our main findings with the simple action-selection model. While it is still possible that the small minority of SNr neurons with sign-reversing responses could mediate antagonistic action selection in the SNr, the majority of SNr neurons could play a fundamentally different function during movement.

Other potential roles for SNr output neurons

If most SNr neurons do not contribute to an antagonistic circuit for action selection in the SNr, what is their function? Thus far, we have emphasized that the *sign* of neural modulation was often congruent between movement types, but the *amplitude* of that modulation often differed. In fact, the majority of SNr neurons exhibited significantly different firing rates for forepaw reaching vs. licking (although surprisingly not for medial vs. lateral reaches made with the same forepaw). Presumably, downstream targets of the SNr could similarly decode movement type from SNr output. It is possible that these signals could then be harnessed for an antagonistic action-selection mechanism, but that antagonism would be played out downstream from the SNr, for example in the brainstem (Lee & Sabatini, 2021).

Notwithstanding, more generally, what other role(s) could be played by the majority of SNr neurons with congruent modulation between actions? It is possible that these neurons *always* decrease (or increase) their activity before *any* movement. In this view, the decrease neurons could nonspecifically disinhibit or gate downstream targets before any movement, with the specific movement dictated by brain structure(s) other than the SNr. Likewise, increase neurons could provide nonspecific inhibition for purposes other than action selection. For example, it has been suggested that cortical-basal ganglia-thalamic circuits act as large-scale feedback loops (Alexander et al., 1986); if so, increase neurons could play a role in stabilizing network output, similar to how negative feedback is commonly used in amplifier circuits to control and linearize output (Sanger, 2003). Feedback inhibition might also provide negative reafference to allow sensory systems to distinguish stimuli arising from external or self-derived sources (von Holst,

1954). While these ideas are of course speculative, they underscore that the common "textbook" assumption that the basal ganglia play a role in real-time action selection is equally speculative – and not well supported by our present findings.

Timing of peri-movement modulation of SNr neurons

We found that pre-movement modulation in many SNr neurons preceded the actual self-timed movement by at least hundreds of ms, in the form of gradual ramping increases or decreases when averaged among trials for single units. The pre-movement activity preceded the earliest EMG activity and could not be attributed to an accelerating pattern of uninstructed jittery movements. We also found early pre-movement modulation in SNr neurons in a "spontaneous" movement task, with no explicit timing requirement. Previous work in monkeys had reported that GPi neurons were modulated late compared to the onset of movement, but those studies generally examined movements made as rapid reactions to external sensory cues; the late timing may have been an idiosyncrasy of rapid reactions. Using self-timed movements, we and others have observed modulation of basal ganglia activity well in advance of movement or EMG activity, in monkey putamen (Lee & Assad, 2003), mouse SNc (Hamilos et al., 2021), mouse striatum (Yang et al., 2024) and now mouse SNr. Thus, for self-timed movements, basal ganglia signals are present early enough to potentially influence movement initiation. Movements made in the absence of external triggers could require a gradual "revving" up of movement circuits that include the basal ganglia, perhaps as a component of a recurrent circuit with thalamus and cortex.

We did not examine the necessity of the SNr for self-timed or spontaneous movements. Lesions of pallidal output nuclei have little effect on reaction times for *cued* movements (Turner & Desmurget, 2010). However, our lab previously found that reversible optogenetic activation or inactivation of SNc dopamine neurons shifted the distribution of self-timed movements to earlier or later times, respectively (Hamilos et al., 2021). Given that the basal ganglia are the major recipient of dopaminergic output from the SNc, perturbations of the SNr may indeed affect self-timed or spontaneous movements. This will be important to examine in future experiments but may require methods to target specific functional cell types in the SNr.

Functional vs. anatomical somatotopy

Somatotopic maps have been described in mouse cortex (Hintiryan et al., 2016) and striatum (Lee et al., 2020), and somatotopy has been inferred in the SNr by tracing monosynaptic striatonigral projections (Foster et al., 2021). However, we did not find evidence for *functional* somatotopy within the SNr in the context of lick vs. reach task, despite recording across the orofacial and forearm regions proposed by anatomical tracing. Rather, we observed a salt-and-pepper mixture of neurons that increased or decreased activity for licking or reaching, across a broad extent of SNr. Sign-reversing neurons, which may specify particular movements, were likewise broadly distributed. Even lick-entrained SNr neurons were not limited to the “lick region” suggested by anatomical tracing. The lack of clear functional somatotopy could arise from input to the SNr from the subthalamic nucleus (STN). STN input to the SNr/GPi is more spatially divergent than striatal input (Mink, 1996), and a recent study reported that STN neurons

are more broadly tuned for multiple movement types (Wu et al., 2025). Regardless, the lack of clear functional organization might distract from more interesting organization at the level of genetically defined cell types. Single-cell RNA sequencing has revealed distinct cellular sub-populations within mouse SNr (Mendelsohn et al., 2024) and entopeduncular nucleus, the rodent homolog of the primate GPi (Wallace et al.), with different projection targets. An intriguing possibility is that the different neuronal responses that we observed in the SNr (i.e., sign-reversing, congruent increase and congruent decrease) could correspond to genetically distinct cell types. Understanding the inputs to and outputs from these neurons could help clarify their functional role. For example, do particular functionally-defined neurons project differentially to brainstem or thalamus/cortex, and can they be distinguished based on patterns of input from the direct or indirect basal ganglia pathways? Addressing these questions will be essential to understanding the role of the basal ganglia output during normal behavior as well as in movement disorders.

Methods

Behavior tasks (self-timed reach/lick, directional reach, and spontaneous reach)

A total of 34 animals (C57BL/6J; 15 female, 19 male; aged 2-12 months) were used in this study. All experiments and protocols were approved by the Harvard Institutional Animal Care and Use Committee and were conducted in accordance with the National Institutes of Health Guide for the Care and Use of Laboratory Animals. Animals were water-restricted 3 days prior to the first training session. Animals received ~1mL of

water or juice each day, adjusted on a daily basis to maintain 80% of their normal bodyweight.

For self-timed reaching (or self-timed licking) tasks, animals were head-fixed and placed on a raised metal platform. An aluminum touch bar was placed 0.5-1 cm in front of the animal, level with or slightly higher than the platform. The touch-bar and lick-spout were each connected to an electronic touch-detection circuit. Before start of a self-timed reach (or lick) trial, the touch-bar (or lick-spout) was deployed into range with a servo motor. When the touch-bar (or lick-spout) was fully deployed (800 ms traversal time), an audio cue (6,272 Hz tone; 100 ms duration) was played to indicate the start of the timing interval. If the first touch (or lick) was made $\geqslant$4 s after cue (correct trial), juice reward was dispensed via a solenoid valve connected to the lick spout; for reach trials, the lick spout was only deployed into licking range when a correct reach was made, effectively delaying reward delivery by 800 ms. After reward delivery, the touch bar and lick spout remains deployed until the end of trial (10 s after cue), at which time they were both retracted via servomotor. In rare cases where animals made a correct movement close to the end of trial, an extra 2 s were allowed for reward collection before retracting the bar and spout. If the first touch (or lick) was made <4 s after cue (incorrect trial), the touch bar (or lick spout) was immediately retracted, followed by a time-out period until end of trial (10 s after cue). If no movement was made during the 10 s trial period, the touch bar (or lick spout) was retracted at 10 s. In all cases (correct, incorrect, no-move) trials were 10 s in length, followed by a random intertrial interval (ITI) ranging between 3-10 s.

For the spontaneous reach task, a small touch-bar (0.4 cm wide, 0.8 cm tall) was placed in front of the contralateral paw. Touching the bar always resulted in its immediate retraction, followed by immediate redeployment (i.e. “cycling” the touch bar, 1400 ms total traversal time). The first reach in a session was always rewarded. Juice was delivered via an advancing lick spout. The lick spout was retracted at the end of the reward collection window (3000 ms, or up to 4000 ms if animals were still licking). The rewarded reach was followed by a random timeout period (drawn from exponential distribution, $\lambda = 20$ s; capped at 10 s to keep sessions reasonably short), during which reaches were not rewarded, and only resulted in “cycling” the touch bar. The first reach after the timeout period was rewarded. No external cue indicated the availability of reward. The exponential time-out period was designed to discourage timing behavior or rhythmic movements. Because many timeout periods were capped at 10 s, the distribution of inter-reach intervals for each behavior session was analyzed to ensure animals made spontaneous movements (exponential distribution) rather than timed movements every 10 seconds (Gaussian distribution).

For the bi-directional reach task, a small aluminum touch-bar (0.4 cm wide, 0.8 cm tall) was presented to the contralateral forepaw (determined relative to SNr recording site). Before the servomotor deployed the target into reaching range, a motorized linear translator first moved the target laterally to one of two pre-defined positions: 0.75 cm lateral to the resting position of the contralateral paw, or 0.5 cm medial to the contralateral paw. Target positions were chosen so that animals were forced to alter their movement trajectories. The target position changed every 15 rewarded trials. The bi-directional reach task is otherwise identical to the spontaneous reach task. (For the

four-directional reach task, two additional targets were added: one in front of the contralateral paw and one in front of the ipsilateral paw. To reduce the complexity of training in the four-directional reach task, animals were rewarded for reaching immediately after the target was deployed. If animals did not move, the target was retracted at 10 seconds and redeployed after a random ITI.)

In all cases, the task flow was controlled in closed-loop with a Teensy microcontroller (PJRC) running a finite state machine program. Behaviorally relevant events (touch-bar contact, lick, cue, reward, etc.) were detected (if not generated) by the microcontroller and sent to the electrophysiology system as digital signals, as well as to a host PC as timestamped events. The host PC ran a custom MATLAB-based graphical user interface, which was used to send control signals (start/stop experiment, change task parameters, etc.), monitor task performance, and log behavioral data to disk.

Surgery and histology

Mice were anaesthetized with isoflurane and head-fixed to a surgical stereotax using earbars. Fur was shaved from the skull, and the scalp disinfected with 70% ethanol and povidone-iodine. Zero pitch was achieved by leveling bregma and lambda points, and zero roll was achieved by leveling two points located 2mm left and right of midline. A lightweight titanium headpost was fixed to the skull with dental cement (C&B Metabond, Parkell). The headpost was lowered and affixed to the skull while held with a custom fitting that ensured the headpost was horizontal in stereotaxic coordinate frame. Craniotomies for chronic or acute recordings were made with a stereotax-mounted drill

and a 0.9 mm diameter drill bit. For chronic recording with microwire bundles, the moveable electrode wires were first covered with petroleum jelly, and the miniature microdrive was cemented to the skull and surrounded by a protective cone and a round cap (made of plastic printer film). For acute recordings, the craniotomy was made under anesthesia 2 hours before the first recording session and covered with a silicone sealant (Kwik-Cast, World Precision Instruments) between sessions.

Stereotactic coordinates for SNr recordings were 1) intermediate-lateral: -3.28 AP/±1.6 ML, used in chronic 32-channel microwire-bundle recording; 2) Intermediate-medial: -3.28 AP/±1.3 ML, for acute 4-shank probes recording, with the most medial shank at ±1.075 ML and the most lateral shank at ±1.525 ML. Depth estimates were made relative to dura, and then converted back to Bregma coordinates; putative SNr units were typically found between -3.8 DV and -5.0 DV.

Histology was performed by euthanizing the animal with pentobarbital sodium and perfusion with 4% PFA (w/v) in 0.1 M PB. Brain tissue was kept in 4% PFA for 24 hours and then in 30% sucrose for at least 24 hours before cutting. Tissue was embedded in OCT (Tissue-Tek) and frozen at -24 °C and sectioned at 50-80 μm thickness. Histology was used to verify implant location of recording electrodes. For chronic recordings, wire bundle tracts were visible in tissue. For acute recordings, the probes were stained with DiI (Thermo Fisher) on the final recording session prior to histology.

Extracellular recording and spike sorting

Initial single-unit recording in SNr was performed with gold-plated NiCr microwires (Sandvik, 12 μm diameter). In early recordings (2 animals) groups of 4 microwires were spun together in tetrode configuration; later recordings used a single bundle of 32 microwires, secured together by ultrathin-walled shrink-wrap (Nordson Medical 103-0402, 0.3 mm internal diameter). The wire tips extended ~1 mm past the end of the tubing, and were allowed to splay out in the brain to sample a larger region. For ground, a single tungsten electrode was placed in the center of the bundle. Electrodes were connected to a 32-channel Omnetics connector via a custom electrode interface board (EIB) and gold pins (NeuraLynx EIB pins). The EIB was mounted on a 3D-printed microdrive headstage, operated by a stainless-steel screw (0.500" length, M1.4-0.3 metric threads). Turning the screw one full turn corresponded to ~300μm change in depth, allowing for granular control of electrode depth after implantation. During training, the tips of the wires were kept ~3.8 mm below dura, roughly 300 μm above the substantia nigra. After the animals were fully trained, we progressively lowered the electrodes into SNr over multiple sessions. The microdrive was moved at least 24 hours before recording sessions to allow wires to fully settle into place. In addition to sampling different depths, moving the electrodes increased the yield compared to a fixed implant.

To better obtain positional information for recorded units, and to improve recording yield, we switched to acute stereotaxic recordings using 128-channel 4-shank silicon probes (model 128DN; Yang et al., 2020). Pictch between shanks was 150 μm, with 32 electrodes arranged 0-775 μm above the tip of each shank. Electrodes on each shank were arranged in staggered columns of two, with 25 μm vertical- and 20 μm horizontal-

pitch. Probes were introduced through a 0.9 mm craniotomy (made on the first day of recording), with the mouse head-fixed using bilateral headpost holders that maintained the headpost (and thus the head) level with respect to the stereotaxic axes.

For acute recordings, two stainless steel ground wires were implanted in the brain and crimped within a single gold-plated connector pin. The craniotomy above the recording site was drilled open before the first recording session. The craniotomy was covered with silicone sealant (Kwik-Cast, World Precision Instruments) between subsequent recording sessions, and bathed in saline during recording. For all recordings, electrodes were gold plated to have 300-750 kΩ impedance.

Signals from extracellular recordings were collected using a Cerebus System (Blackrock Microsystems) or an RHD2000 System (Intan Technologies). Amplified electrode signals was sampled at 30,000 Hz and band-pass filtered (250Hz-7500Hz Butterworth) before storing to disk. Relevant behavior events were recorded as digital signals and stored as timestamps.

Spike-sorting was completed using custom software in MATLAB. The threshold for spike-detection was set as 2.5 SD of the signal (SD is estimated from MAD/0.6745): this usually captured spike waveforms as well as noise. Spike waveforms were separated from noise via k-means clustering, or by fitting a Gaussian mixture model, then visualized after dimensionality-reduction (from 31D to 3D) via principal component analysis. Because simple unsupervised clustering methods sometimes produced poor separation, we improved automated results by manual curation (removing electrical artefacts, splitting, merging, and adjusting cluster boundaries). We aimed to extract only high signal-to-noise single-unit spike trains from recorded data. Units with low signal-to-

noise ratios were considered multi-unit activity and not included in analysis. Additionally, inter-spike intervals (ISI) smaller than 1.5 ms were at odds with the refractory period of spiking neurons. Therefore, units were classified as *single units* only if no more than 5% of interspike intervals were shorter than 1.5 ms. In cases where the same neuron was recorded on multiple adjacent channels, duplicate units were detected via cross-correlation analysis and removed. All presented neuronal data were from well-isolated single units.

In particular, we checked carefully that spike waveforms of “single units” remained uniform throughout self-timed trials, to ensure that any observed *changes* in firing rates were not due to spike-sorting artifacts, other units interfering, etc. In some cases electrode drift caused units to “disappear” halfway through recording sessions: these drifting units were excluded from analysis.

SNr neurons are tonically active, with typical spiking rates ranging from 20 to 100 spikes per second (sp/s). Dopamine neurons in the SNc and dorsal parts of SNr stereotypically exhibit 6-8 sp/s baseline firing and wider spike waveforms. Therefore, only single units with baseline firing rates >15 sp/s were considered putative SNr neurons.

Identifying the sign and onset of pre-movement neural modulation

For pre-movement modulation in self-timed (reaching or licking) tasks, only trials in which the animal waited $\geqslant$2 s after cue were included, since shorter trials may include reactionary movements to the cue. A peri-event time histogram (PETH) was generated

for each SNr unit by counting spikes in 100-ms (or 25-ms, for visualization) bins aligned to movement time (bar or spout contact for reach or lick tasks, respectively). Spike counts are divided by bin width and then averaged across trials to estimate spike rates. For comparing across multiple units, the PETH for each unit is normalized by subtracting the mean and dividing by the standard deviation estimated during the baseline window ([-4000 ms, -2000 ms] prior to movement). For calculating response onset latency, a smoothed PETH was generated by first counting spikes using 1-ms bins, then convolving with a Gaussian smoothing kernel (σ=100 ms, width=1000 ms), then averaging across trials.

In order to test whether putative SNr neurons were significantly modulated by movement, on a unit-by-unit basis, we first calculated its PETH (100-ms bins) aligned to the first movement (bar/spout contact) in each self-timed trial. We then compiled the distribution of its spiking rate in the *peri-movement window* ([-300, 0] ms before touchbar contact), as well as the distribution of spiking rates in the *baseline window* ([-4000, -2000] ms before touchbar contact). This yielded a distribution of peri-movement spiking rates (sample size $= 3K$, where K is the number of trials, and 3 is the number of 100-ms bins in the peri-movement window), as well as a distribution of baseline spiking rates (sample size $= 20K$, across 20 100-m bins). We then measured the difference between the means of the peri-movement vs. baseline distributions, using a permutation test to assess whether this difference is statistically significant, under the null hypothesis that spike rates in each of the $20K$ baseline bins and $3K$ pre-movement bins are independent and identically distributed (i.i.d.). Spike rates in baseline and pre-movement bins were shuffled 100,000 times. For each permutation, we re-sampled $20K$

‘baseline’ bins and $3K$ ‘response’ bins and calculated the difference between their means. The 99% confidence interval for this difference-of-means was calculated from 100,000 permutations. A unit was classified as significantly movement-responsive if the observed difference-of-means fell outside the 99% confidence interval (i.e., permutation test; $p<0.01$).

For responsive SNr units, we then analyzed their response latency (i.e., onset time) using the kernel-smoothed PETH previously described. Looking back from movement time, response onset was defined as when the normalized, smoothed spike rates (sampled at 1 ms resolution) first exhibited 25 ms of baseline activity (<0.25 SD) followed by 50 ms of above-baseline (or below-baseline) activity ($\geqslant$0.25 SD).

Video tracking and movement onset detection

Bilateral video recordings of animal movement were collected using webcams at ~30 frames per second. IR LEDs are used to illuminate the chamber during recording. Video streaming and logging was controlled via MATLAB (Image Acquisition Toolbox). A timestamp is registered by MATLAB on every 10th frame, which is used to align video and spiking data. DeepLapCut (Mathis et al., 2018) was used to train a neural network (ResNet50) to identify the location of relevant body parts (front and rear paws, nose, tail, and the midpoint of the spine). Model-predicted positions with low likelihoods (<95%) were discarded, because these were typically erroneous predictions due to occlusion of body parts.

To compare movement onset times against neural onset times, video frames were aligned with electrophysiology data (made possible by saving the timestamp of every 10th video frame, as well as the timestamp of the first electrophysiology sample). Accurate timing alignment was verified by having a LED visible to the camera turn on at the start of each trial, while the LED gating signal was recorded as a digital input on the spike-acquisition signal.

Electrically detected touchbar contact times were used to determine the end of arm reaches. Movement onset (forepaw lift) time was estimated from video-derived trajectories of the contralateral forepaw (resampled at 25-ms intervals), by tracing backwards from end-of-reach until the contralateral forepaw was observed at its resting position for 25 ms followed by 50 ms away from resting position ($\geqslant$0.25 SD). Automatically detected movement-onset times as well as bar-contact times have been validated by manual inspection of sample trials.

In one animal (3 sessions), a pair of EMG wires were implanted in the neck muscles. The EMG signal were amplified (x1000 gain) and filtered (50-500 Hz band-pass) using an amplifier (A-M Systems, Model 1700), half-wave rectified and recorded at 30,000 Hz (Intan RHD2000 evaluation board), then smoothed (33.3ms moving-average) and normalized (robust z-score) in MATLAB. To detect EMG onset, the smoothed, normalized EMG signal was resampled at 25-ms intervals, and EMG onset was determined by tracing backwards to the first time EMG exceeded 0.25 SD for two consecutive 25-ms samples.

Generalized linear model for decoding SNr unit activity

For each SNr units, a series of nested GLMs were fitted to predict its spike rate time course. 4 categories of predictors were iteratively added to the model, with the goal of estimating whether model performance had been improved by including more predictors:

1) The *cue* predictor was generated by converting discrete trial-start events into a continuous time series across the session, using a set of 13 overlapping cosine kernels (width 400 ms, temporal offsets ranging from 0 to ±600 ms). Negative latency offsets were included because the servo motors engage 800 ms before trial start, a potentially salient cue for behaving animals.

2) The *velocity* predictors were generated from video analysis of body parts. Position traces detected by neuron networks were converted into 2D velocity traces, then normalized by the mean and SD (i.e., z-scoring). For bilateral videos, the x and y axes respectively point to anterior and dorsal side of the animal. Body parts visible to both cameras (nose, spine, tail) were averaged between cameras after normalization. Normalized velocity signals were then convolved with a series of 9 cosine kernels (width 100 ms, temporal offsets 0 to ±200 ms). Latency offsets of up to 200ms covered the range of neural response times previously observed in monkey SNr and GPi.

3) The *ramp* predictors were a series of linear ramps (0 to 1) terminating at touchbar-contact times; 19 such ramps starting 100-2000 ms prior to touchbar contact were added to the model. They were timing-invariant in the sense that these ramps did not begin at cue (when the animal begins timing) but were rather aligned to the movement time. We added the ramp predictors *after* the velocity predictors in the nested GLM

approach: if the velocity predictors already explained the ramping activity, then adding the ramp predictors should not significantly improve model performance.

4) The *trial-progression* predictors were signals that increased linearly from 0 to 1 from trial-start to touchbar contact. These predictors capture elapsed time from cue up to the moment of movement.

GLMs were fit using iteratively reweighted least squares to find the maximum likelihood estimate for model parameters (Statistics and Machine Learning Toolbox, MATLAB). Five-fold cross validation was employed to reduce overfitting. Model performance was evaluated by variance explained (R^2) and Akaike Information Criterion (AIC), calculating the additional variance explained (ΔR^2) as well as the reduction of Akaike Information Criterion (ΔAIC) after adding more predictors to the model. Fitted models were used to predict single-trial spiking rate traces, which were then averaged across trials and compared to true observations.

Linear discriminant analysis for decoding SNr population activity

Population SNr activity was collected during the peri-movement period for reach vs. lick trials ([-0.3, 0] s relative to bar/spout contact) and during the two-target lateral vs. medial reach trials ([-0.1, 0.2] s relative to reach onset). To ensure population size, only sessions with $\geqslant$10 recorded SNr units were included (10 animals/13 session for reach vs. lick; 3 animals/7 sessions for lateral vs. medial reach). Peri-movement spiking activity for each SNr unit was normalized to the mean and SD of its baseline spike rates (estimated during baseline window: [-4, -2] s before movement). For each session,

normalized peri-movement activity as well as baseline activity were used to train a decoding model (linear discriminant analysis, or LDA) by fitting a Gaussian mixture model (assuming equal variance for all 3 states: lick/reach/baseline; or lateral-reach/medial-reach/baseline). We assumed empirical prior probabilities proportional to the number of observed trials of each type (490 reach vs. 680 lick trials vs. 1170 baseline periods; 351 lateral reach trials vs. 424 medial reach trials vs. 775 baseline periods). We implemented five-fold cross-validation to reduce overfitting. Models trained on the peri-movement window ([-0.3, 0] s or [-0.1, 0.2] s) were used to decode population neural activity from [-4, 2] s relative to movement; yielding model-predicted probabilities for each state: p(reach), p(lick) and p(baseline); or p(lateral), p(medial) and p(baseline) for the reach vs. lick task and two-target reach task, respectively. For each time point t (relative to movement time), we then took the difference in predicted probabilities: $\Delta p(t) = p(\text{reach}|t) - p(\text{lick}|t)$, or $\Delta p(t) = p(\text{lateral}|t) - p(\text{medial}|t)$ as a one-dimensional representation of the time-course of SNr population activity. Model predictions are averaged across true reach trials vs. true lick trials (or true lateral-reach trials vs. true medial-reach trials). For visualization, we then averaged decoded population activity across sessions (weighted by trial count). This grand average is representative of each individual session upon inspection.

To examine the significance of model-decoded activity, we also trained null models using shuffled training data where the lick/reach (or lateral/medial reach) trial labels were randomly permuted. 10,000 permutations were performed, each time generating a null model used to decode population activity. As expected, the null models were not

able to decode SNr activity. The 99% confidence interval (i.e., 0.5% and 99.5% percentiles) for $\Delta p(t)$ was calculated amongst the 10,000 null models.

Detecting differences across PETHs grouped by movement time

Self-timed reach trials were grouped into 4 bins according to movement time (2-4s, 4-6s, 6-8s, and 8-10s; i.e., how long animals waited after the start cue before touching the bar; trials shorter than 2s were considered acute reactions to the cue and excluded from analysis). Thus, for each SNr unit, 4 peri-reach spike-time histograms (PETH) were generated for the 4 groups of trials. 6 pairwise Euclidean distances were calculated for the 4 PETHs. Distances were only calculated in the window [-2, 0] s before bar contact where the majority of pre-movement activity was observed. We took d, the mean across the 6 Euclidean distances, as a measurement of the dissimilarity between ramping neural activity across different movement times. Under the null hypothesis that there is no difference between groups ($d = 0$), we randomly shuffled trial identities before grouping and calculating the 4 PETHs, and generated a 99% confidence interval for d. The minority (30/641) of SNr units whose d fell outside the confidence interval were potentially timing-dependent and selected for further inspection.

Detection of lick-entrained units using circular statistics

We used circular statistics to determine if the spiking rates of SNr units fluctuated in sync with rhythmic lick bouts. The interval between two adjacent licks (spout-contact)

were evenly divided into 30 bins to generate a circular PETH, where each bin in the PETH represents the average spike rate during a certain phase (0-2π) of the lick. Interlick intervals longer than 200 ms (non-rhythmic licking) or shorter than 50 ms (likely due to lick-detector artifact) were not considered. For each unit, its circular PETH was expressed as 30 vectors in polar coordinates: (phase, spike rate). Averaging across these 30 vectors in Euclidean space yielded a “phase-selectivity vector”, whose magnitude should be zero for units with perfectly stationary spike rates. A non-zero magnitude indicated increased or decreased spike rates during certain phases of licking (i.e., lick-entrained). Therefore, we tested whether the average vector had a significantly non-zero magnitude using a permutation test as follows: 1) shuffle the phase of all 30 PETH vectors 10,000 times, each time calculating the magnitude of the average vector; 2) check whether the “observed” average vector magnitude falls within the top 1% of the distribution of “shuffled” vector magnitudes (i.e., $p < 0.01$).

While lick-entrainment was detected using lick-pairs, we tried to verify whether entrainment was present in longer lick bouts. To that end we visualized lick-entrainment by plotting an extended PETH spanning 5 consecutive licks. Longer lick trains were split into five-lick bouts (whereas shorter lick bouts were discarded); all five-lick bouts in a session were aligned by phase (0-8π, with multiples of 2π denoting spout contact) to generate a PETH (as seen in Figure 6b&d) showing lick-entrainment.

References

Aldridge, J. W., Anderson, R. J., & Murphy, J. T. (1980). The role of the basal ganglia in controlling a movement initiated by a visually presented cue. *Brain Research*, *192*(1), 3-16. https://doi.org/https://doi.org/10.1016/0006-8993(80)91003-3

Alexander, G. E., & Crutcher, M. D. (1990). Functional architecture of basal ganglia circuits: neural substrates of parallel processing. *Trends in Neurosciences*, *13*(7), 266-271.

Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu Rev Neurosci*, *9*, 357-381. https://doi.org/10.1146/annurev.ne.09.030186.002041

Anderson, M. E., & Horak, F. B. (1985). Influence of the globus pallidus on arm movements in monkeys. III. Timing of movement-related information. *Journal of neurophysiology*, *54*(2), 433-448. https://doi.org/10.1152/jn.1985.54.2.433

Arber, S., & Costa, R. M. (2022). Networking brainstem and basal ganglia circuits for movement. *Nature Reviews Neuroscience*, *23*(6), 342-360. https://doi.org/10.1038/s41583-022-00581-w

Bariselli, S., Fobbs, W. C., Creed, M. C., & Kravitz, A. V. (2019). A competitive model for striatal action selection. *Brain Research*, *1713*, 70-79. https://doi.org/https://doi.org/10.1016/j.brainres.2018.10.009

Brotchie, P., Iansek, R., & Horne, M. K. (1991). Motor function of the monkey globus pallidus. 2. Cognitive aspects of movement and phasic neuronal activity. *Brain*, *114 (Pt 4)*, 1685-1702. https://doi.org/10.1093/brain/114.4.1685

Bryden, D. W., Johnson, E. E., Diao, X., & Roesch, M. R. (2011). Impact of expected value on neural activity in rat substantia nigra pars reticulata. *European Journal of Neuroscience*, *33*(12), 2308-2317. https://doi.org/https://doi.org/10.1111/j.1460-9568.2011.07705.x

Chen, Z., Zhang, Z.-Y., Zhang, W., Xie, T., Li, Y., Xu, X.-H., & Yao, H. (2021). Direct and indirect pathway neurons in ventrolateral striatum differentially regulate licking movement and nigral responses. *Cell Reports*, *37*(3). https://doi.org/10.1016/j.celrep.2021.109847

Cruz, B. F., Guiomar, G., Soares, S., Motiwala, A., Machens, C. K., & Paton, J. J. (2022). Action suppression reveals opponent parallel control via striatal circuits. *Nature*, *607*(7919), 521-526. https://doi.org/10.1038/s41586-022-04894-9

Cui, G., Jun, S. B., Jin, X., Pham, M. D., Vogel, S. S., Lovinger, D. M., & Costa, R. M. (2013). Concurrent activation of striatal direct and indirect pathways during action initiation. *Nature*, *494*(7436), 238-242. https://doi.org/10.1038/nature11846

Deniau, J. M., & Chevalier, G. (1992). The lamellar organization of the rat substantia nigra pars reticulata: Distribution of projection neurons. *Neuroscience*, *46*(2), 361-377. https://doi.org/https://doi.org/10.1016/0306-4522(92)90058-A

Deniau, J. M., Mailly, P., Maurice, N., & Charpier, S. (2007). The pars reticulata of the substantia nigra: a window to basal ganglia output. In J. M. Tepper, E. D. Abercrombie, & J. P. Bolam (Eds.), *Progress in brain research* (Vol. 160, pp.

151-172). Elsevier. https://doi.org/https://doi.org/10.1016/S0079-6123(06)60009-5

Desmurget, M., & Turner, R. S. (2008). Testing Basal Ganglia Motor Functions Through Reversible Inactivations in the Posterior Internal Globus Pallidus. *Journal of neurophysiology*, *99*(3), 1057-1076. https://doi.org/10.1152/jn.01010.2007

Fan, D., Rossi, M. A., & Yin, H. H. (2012). Mechanisms of Action Selection and Timing in Substantia Nigra Neurons. *The Journal of Neuroscience*, *32*(16), 5534-5548. https://doi.org/10.1523/jneurosci.5924-11.2012

Foster, N. N., Barry, J., Korobkova, L., Garcia, L., Gao, L., Becerra, M., Sherafat, Y., Peng, B., Li, X., Choi, J.-H., Gou, L., Zingg, B., Azam, S., Lo, D., Khanjani, N., Zhang, B., Stanis, J., Bowman, I., Cotter, K., . . . Dong, H.-W. (2021). The mouse cortico–basal ganglia–thalamic network. *Nature*, *598*(7879), 188-194. https://doi.org/10.1038/s41586-021-03993-3

Frost-Nylén, J., Thompson, W. S., Robertson, B., & Grillner, S. (2024). The Basal Ganglia Downstream Control of Action – An Evolutionarily Conserved Strategy. *Current Neuropharmacology*, *22*(9), 1419-1430. https://doi.org/10.2174/1570159X21666230810141746

Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*, *107*(2), 289-344. https://doi.org/10.1037/0033-295X.107.2.289

Georgopoulos, A., DeLong, M., & Crutcher, M. (1983). Relations between parameters of step-tracking movements and single cell discharge in the globus pallidus and subthalamic nucleus of the behaving monkey. *The Journal of Neuroscience*, *3*(8), 1586-1598. https://doi.org/10.1523/jneurosci.03-08-01586.1983

Gerfen, C. R. (2023). Segregation of D1 and D2 dopamine receptors in the striatal direct and indirect pathways: An historical perspective [Review]. *Frontiers in Synaptic Neuroscience*, *14*. https://doi.org/10.3389/fnsyn.2022.1002960

Glickstein, M., & Stein, J. (1991). Paradoxical movement in Parkinson's disease. *Trends in Neurosciences*, *14*(11), 480-482. https://doi.org/https://doi.org/10.1016/0166-2236(91)90055-Y

Gulley, J. M., Kosobud, A. E. K., & Rebec, G. V. (2002). Behavior-related modulation of substantia nigra pars reticulata neurons in rats performing a conditioned reinforcement task. *Neuroscience*, *111*(2), 337-349. https://doi.org/https://doi.org/10.1016/S0306-4522(02)00018-0

Hamilos, A. E., Spedicato, G., Hong, Y., Sun, F., Li, Y., & Assad, J. A. (2021). Slowly evolving dopaminergic activity modulates the moment-to-moment probability of reward-related self-timed movements. *eLife*, *10*, e62583. https://doi.org/10.7554/eLife.62583

Handel, A., & Glimcher, P. W. (1999). Quantitative Analysis of Substantia Nigra Pars Reticulata Activity During a Visually Guided Saccade Task. *Journal of neurophysiology*, *82*(6), 3458-3475. https://doi.org/10.1152/jn.1999.82.6.3458

Hasnain, M. A., Birnbaum, J. E., Ugarte Nunez, J. L., Hartman, E. K., Chandrasekaran, C., & Economo, M. N. (2025). Separating cognitive and motor processes in the behaving mouse. *Nature Neuroscience*. https://doi.org/10.1038/s41593-024-01859-1

Hikosaka, O., & Wurtz, R. H. (1983a). Visual and oculomotor functions of monkey substantia nigra pars reticulata. I. Relation of visual and auditory responses to

saccades. *Journal of neurophysiology*, *49*(5), 1230-1253. https://doi.org/10.1152/jn.1983.49.5.1230

Hikosaka, O., & Wurtz, R. H. (1983b). Visual and oculomotor functions of monkey substantia nigra pars reticulata. II. Visual responses related to fixation of gaze. *Journal of neurophysiology*, *49*(5), 1254-1267. https://doi.org/10.1152/jn.1983.49.5.1254

Hikosaka, O., & Wurtz, R. H. (1983c). Visual and oculomotor functions of monkey substantia nigra pars reticulata. IV. Relation of substantia nigra to superior colliculus. *Journal of neurophysiology*, *49*(5), 1285-1301. https://doi.org/10.1152/jn.1983.49.5.1285

Hintiryan, H., Foster, N. N., Bowman, I., Bay, M., Song, M. Y., Gou, L., Yamashita, S., Bienkowski, M. S., Zingg, B., Zhu, M., Yang, X. W., Shih, J. C., Toga, A. W., & Dong, H.-W. (2016). The mouse cortico-striatal projectome. *Nature Neuroscience*, *19*(8), 1100-1114. https://doi.org/10.1038/nn.4332

Horak, F. B., & Anderson, M. E. (1984). Influence of globus pallidus on arm movements in monkeys. I. Effects of kainic acid-induced lesions. *Journal of neurophysiology*, *52*(2), 290-304. https://doi.org/10.1152/jn.1984.52.2.290

Joseph, J. P., & Boussaoud, D. (1985). Role of the cat substantia nigra pars reticulata in eye and head movements I. Neural activity. *Experimental Brain Research*, *57*(2), 286-296. https://doi.org/10.1007/BF00236534

Joseph, J. P., Boussaoud, D., & Biguer, B. (1985). Activity of neurons in the cat substantia nigra pars reticulata during drinking. *Experimental Brain Research*, *60*(2), 375-379. https://doi.org/10.1007/BF00235932

Kaneda, K., Isa, K., Yanagawa, Y., & Isa, T. (2008). Nigral Inhibition of GABAergic Neurons in Mouse Superior Colliculus. *The Journal of Neuroscience*, *28*(43), 11071-11078. https://doi.org/10.1523/jneurosci.3263-08.2008

Klaus, A., Alves da Silva, J., & Costa, R. M. (2019). What, If, and When to Move: Basal Ganglia Circuits and Self-Paced Action Initiation. *Annual Review of Neuroscience*, *42*(Volume 42, 2019), 459-483. https://doi.org/https://doi.org/10.1146/annurev-neuro-072116-031033

Lee, I. H., & Assad, J. A. (2003). Putaminal Activity for Simple Reactions or Self-Timed Movements. *Journal of neurophysiology*, *89*(5), 2528-2537. https://doi.org/10.1152/jn.01055.2002

Lee, I. H., Seitz, A. R., & Assad, J. A. (2006). Activity of Tonically Active Neurons in the Monkey Putamen During Initiation and Withholding of Movement. *Journal of neurophysiology*, *95*(4), 2391-2403. https://doi.org/10.1152/jn.01053.2005

Lee, J., & Sabatini, B. L. (2021). Striatal indirect pathway mediates exploration via collicular competition. *Nature*, *599*(7886), 645-649. https://doi.org/10.1038/s41586-021-04055-4

Lee, J., Wang, W., & Sabatini, B. L. (2020). Anatomically segregated basal ganglia pathways allow parallel behavioral modulation. *Nature Neuroscience*, *23*(11), 1388-1398. https://doi.org/10.1038/s41593-020-00712-5

Lintz, M. J., & Felsen, G. (2016). Basal ganglia output reflects internally-specified movements. *eLife*, *5*, e13833. https://doi.org/10.7554/eLife.13833

Magariños-Ascone, C., Buño, W., & García-Austt, E. (1992). Activity in monkey substantia nigra neurons related to a simple learned movement. *Experimental Brain Research*, *88*(2), 283-291. https://doi.org/10.1007/BF02259103

Maimon, G., & Assad, J. A. (2006). A cognitive signal for the proactive timing of action in macaque LIP. *Nature Neuroscience*, *9*(7), 948-955. https://doi.org/10.1038/nn1716

Markowitz, J. E., Gillis, W. F., Beron, C. C., Neufeld, S. Q., Robertson, K., Bhagat, N. D., Peterson, R. E., Peterson, E., Hyun, M., Linderman, S. W., Sabatini, B. L., & Datta, S. R. (2018). The Striatum Organizes 3D Behavior via Moment-to-Moment Action Selection. *Cell*, *174*(1), 44-58.e17. https://doi.org/10.1016/j.cell.2018.04.019

Martin, J. P. (1967). *The Basal Ganglia and Posture*. Pitman Medical. https://books.google.com/books?id=lFU8XwAACAAJ

Mathis, A., Mamidanna, P., Cury, K. M., Abe, T., Murthy, V. N., Mathis, M. W., & Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience*, *21*(9), 1281-1289. https://doi.org/10.1038/s41593-018-0209-y

McElvain, L. E., Chen, Y., Moore, J. D., Brigidi, G. S., Bloodgood, B. L., Lim, B. K., Costa, R. M., & Kleinfeld, D. (2021). Specific populations of basal ganglia output neurons target distinct brain stem areas while collateralizing throughout the diencephalon. *Neuron*, *109*(10), 1721-1738.e1724. https://doi.org/10.1016/j.neuron.2021.03.017

Mendelsohn, A. I., Nikoobakht, L., Bikoff, J. B., & Costa, R. M. (2024). Segregated basal ganglia output pathways correspond to genetically divergent neuronal subclasses. *bioRxiv*. https://doi.org/10.1101/2024.08.28.610136

Mink, J. W. (1996). The basal ganglia: focused selection and inhibition of competing motor programs. *Prog Neurobiol*, *50*(4), 381-425. https://doi.org/10.1016/s0301-0082(96)00042-1

Mink, J. W., & Thach, W. T. (1991a). Basal ganglia motor control. II. Late pallidal timing relative to movement onset and inconsistent pallidal coding of movement parameters. *Journal of neurophysiology*, *65*(2), 301-329. https://doi.org/10.1152/jn.1991.65.2.301

Mink, J. W., & Thach, W. T. (1991b). Basal ganglia motor control. III. Pallidal ablation: normal reaction time, muscle cocontraction, and slow movement. *Journal of neurophysiology*, *65*(2), 330-351. https://doi.org/10.1152/jn.1991.65.2.330

Mink, J. W., & Thach, W. T. (1993). Basal ganglia intrinsic circuits and their role in behavior. *Current Opinion in Neurobiology*, *3*(6), 950-957. https://doi.org/https://doi.org/10.1016/0959-4388(93)90167-W

Mitchell, S. J., Richardson, R. T., Baker, F. H., & DeLong, M. R. (1987). The primate globus pallidus: neuronal activity related to direction of movement. *Experimental Brain Research*, *68*(3), 491-505. https://doi.org/10.1007/BF00249793

Mushiake, H., & Strick, P. L. (1995). Pallidal neuron activity during sequential arm movements. *Journal of neurophysiology*, *74*(6), 2754-2758. https://doi.org/10.1152/jn.1995.74.6.2754

Nambu, A. (2007). Globus pallidus internal segment. In J. M. Tepper, E. D. Abercrombie, & J. P. Bolam (Eds.), *Progress in brain research* (Vol. 160, pp.

135-150). Elsevier. https://doi.org/https://doi.org/10.1016/S0079-6123(06)60008-3

Pillow, J. (2007). Likelihood-based approaches to modeling the neural code. *Bayesian brain: Probabilistic approaches to neural coding*, *70*, 53-70.

Rossi, M. A., Li, H. E., Lu, D., Kim, I. H., Bartholomew, R. A., Gaidis, E., Barter, J. W., Kim, N., Cai, M. T., Soderling, S. H., & Yin, H. H. (2016). A GABAergic nigrotectal pathway for coordination of drinking behavior. *Nature Neuroscience*, *19*(5), 742-748. https://doi.org/10.1038/nn.4285

Russo, Y., Stuart, S., Silva-Batista, C., Brumbach, B., Vannozzi, G., & Mancini, M. (2022). Does visual cueing improve gait initiation in people with Parkinson's disease? *Human Movement Science*, *84*, 102970. https://doi.org/https://doi.org/10.1016/j.humov.2022.102970

Sanger, T. D. (2003). Childhood onset generalised dystonia can be modelled by increased gain in the indirect basal ganglia pathway. *Journal of Neurology, Neurosurgery & Psychiatry*, *74*(11), 1509-1515. https://doi.org/10.1136/jnnp.74.11.1509

Sarma, S. V., Cheng, M. L., Eden, U. T., Williams, Z., Brown, E. N., & Eskandar, E. (2012). The effects of cues on neurons in the basal ganglia in Parkinson's disease [Original Research]. *Frontiers in Integrative Neuroscience*, *6*. https://doi.org/10.3389/fnint.2012.00040

Sato, M., & Hikosaka, O. (2002). Role of Primate Substantia Nigra Pars Reticulata in Reward-Oriented Saccadic Eye Movement. *The Journal of Neuroscience*, *22*(6), 2363-2373. https://doi.org/10.1523/jneurosci.22-06-02363.2002

Schroll, H., & Hamker, F. H. (2013). Computational models of basal-ganglia pathway functions: focus on functional neuroanatomy [Review]. *Frontiers in systems neuroscience*, *7*. https://doi.org/10.3389/fnsys.2013.00122

Schultz, W. (1986). Activity of pars reticulata neurons of monkey substantia nigra in relation to motor, sensory, and complex events. *Journal of neurophysiology*, *55*(4), 660-677. https://doi.org/10.1152/jn.1986.55.4.660

Starkweather, C. K., Babayan, B. M., Uchida, N., & Gershman, S. J. (2017). Dopamine reward prediction errors reflect hidden-state inference across time. *Nature Neuroscience*, *20*(4), 581-589. https://doi.org/10.1038/nn.4520

Turner, R. S., & Anderson, M. E. (1997). Pallidal Discharge Related to the Kinematics of Reaching Movements in Two Dimensions. *Journal of neurophysiology*, *77*(3), 1051-1074. https://doi.org/10.1152/jn.1997.77.3.1051

Turner, R. S., & Desmurget, M. (2010). Basal ganglia contributions to motor control: a vigorous tutor. *Current Opinion in Neurobiology*, *20*(6), 704-716. https://doi.org/https://doi.org/10.1016/j.conb.2010.08.022

Villalobos, C. A., & Basso, M. A. (2022). Optogenetic activation of the inhibitory nigro-collicular circuit evokes contralateral orienting movements in mice. *Cell Reports*, *39*(3). https://doi.org/10.1016/j.celrep.2022.110699

von Holst, E. (1954). Relations between the central Nervous System and the peripheral organs. *The British Journal of Animal Behaviour*, *2*(3), 89-94. https://doi.org/https://doi.org/10.1016/S0950-5601(54)80044-X

Wallace, M. L., Saunders, A., Huang, K. W., Philson, A. C., Goldman, M., Macosko, E. Z., McCarroll, S. A., & Sabatini, B. L. (2017). Genetically Distinct Parallel

Pathways in the Entopeduncular Nucleus for Limbic and Sensorimotor Output of the Basal Ganglia. *Neuron*, *94*(1), 138-152.e135. https://doi.org/10.1016/j.neuron.2017.03.017

Wichmann, T., & Kliem, M. A. (2004). Neuronal Activity in the Primate Substantia Nigra Pars Reticulata During the Performance of Simple and Memory-Guided Elbow Movements. *Journal of neurophysiology*, *91*(2), 815-827. https://doi.org/10.1152/jn.01180.2002

Wu, B.-S., Ming, M.-Y., & Wu, Y.-W. (2025). Mixed Selectivity of Subthalamic Nucleus Neurons in Encoding Motor and Reward Behaviors. *bioRxiv*, 2025.2002.2012.637797. https://doi.org/10.1101/2025.02.12.637797

Yang, L., Lee, K., Villagracia, J., & Masmanidis, S. C. (2020). Open source silicon microprobes for high throughput neural recording. *Journal of Neural Engineering*, *17*(1), 016036. https://doi.org/10.1088/1741-2552/ab581a

Yang, Z., Inagaki, M., Gerfen, C. R., Fontolan, L., & Inagaki, H. K. (2024). Integrator dynamics in the cortico-basal ganglia loop underlie flexible motor timing. *bioRxiv*, 2024.2006.2029.601348. https://doi.org/10.1101/2024.06.29.601348

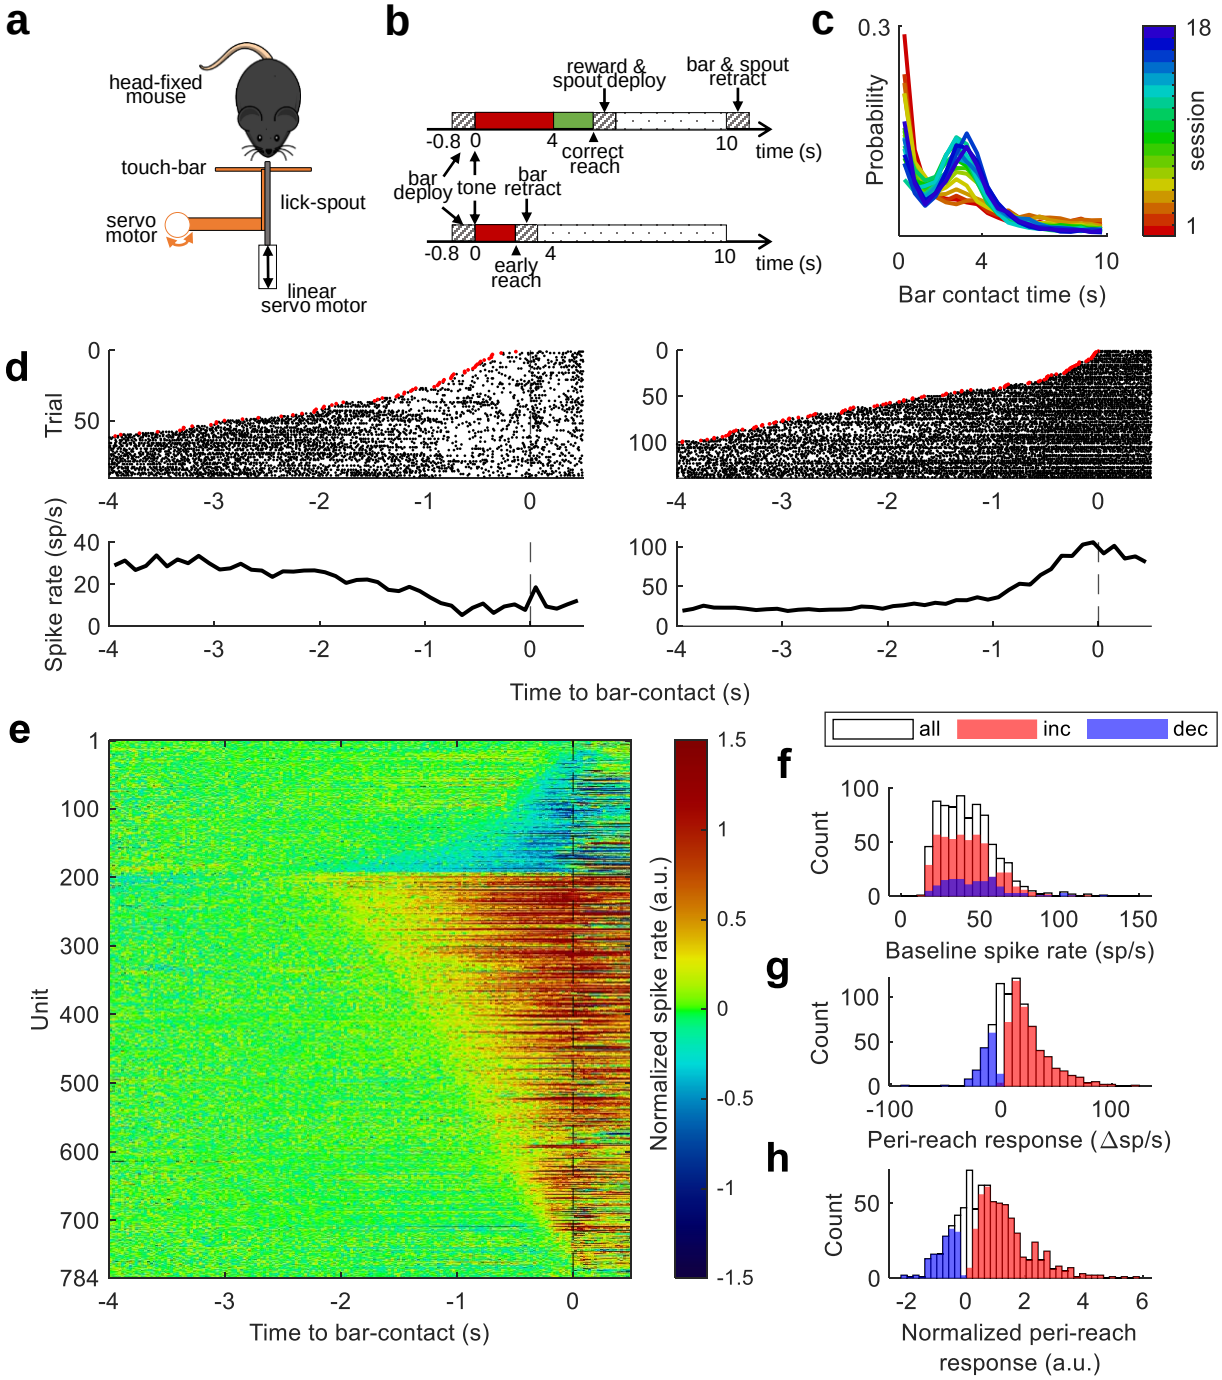

Figure 1. SNr neural activity in the self-timed forepaw reach task

a. Schematic of the self-timed forepaw reach task. **b.** Example time course of a rewarded (top) vs. unrewarded (bottom) trial, contingent on time of first reach. **c.** Distribution of movement times (i.e. first bar-contact after cue) relative to the start cue, shown in different colors for the first 18 sessions (n=29 animals). **d. Top:** spike rasters for two putative SNr neurons in the self-timed reach task, aligned to time of bar-contact; trials are sorted by movement time; spikes are not shown before the start cue (red dot); **bottom:** peri-event time histograms (PETHs) aligned to bar-contact time, for the same two SNr units. **e.** Normalized peri-movement PETHs (25ms binned spike counts) for 784 SNr units in the self-timed reach task, visualized as a heat map. Spike rates are normalized on a per-unit basis, relative to the mean and standard deviation during the "baseline window" ([-4, -2] s before bar-contact). Units are ordered by the sign and onset latency of the peri-reach response. **f.** Distribution of median spike rate during intertrial intervals. **g.** Distribution of peri-movement response, measured as difference between the peri-movement spike rate vs. baseline spike rate ([-0.3, 0] s and [-4, -2] s relative to bar-contact, respectively). **h.** Distribution of peri-movement responses normalized to the mean and S.D. of baseline activity. **f-h:** white: all SNr units (n=784); red: increase units (n=496); blue: decrease units (n=145).

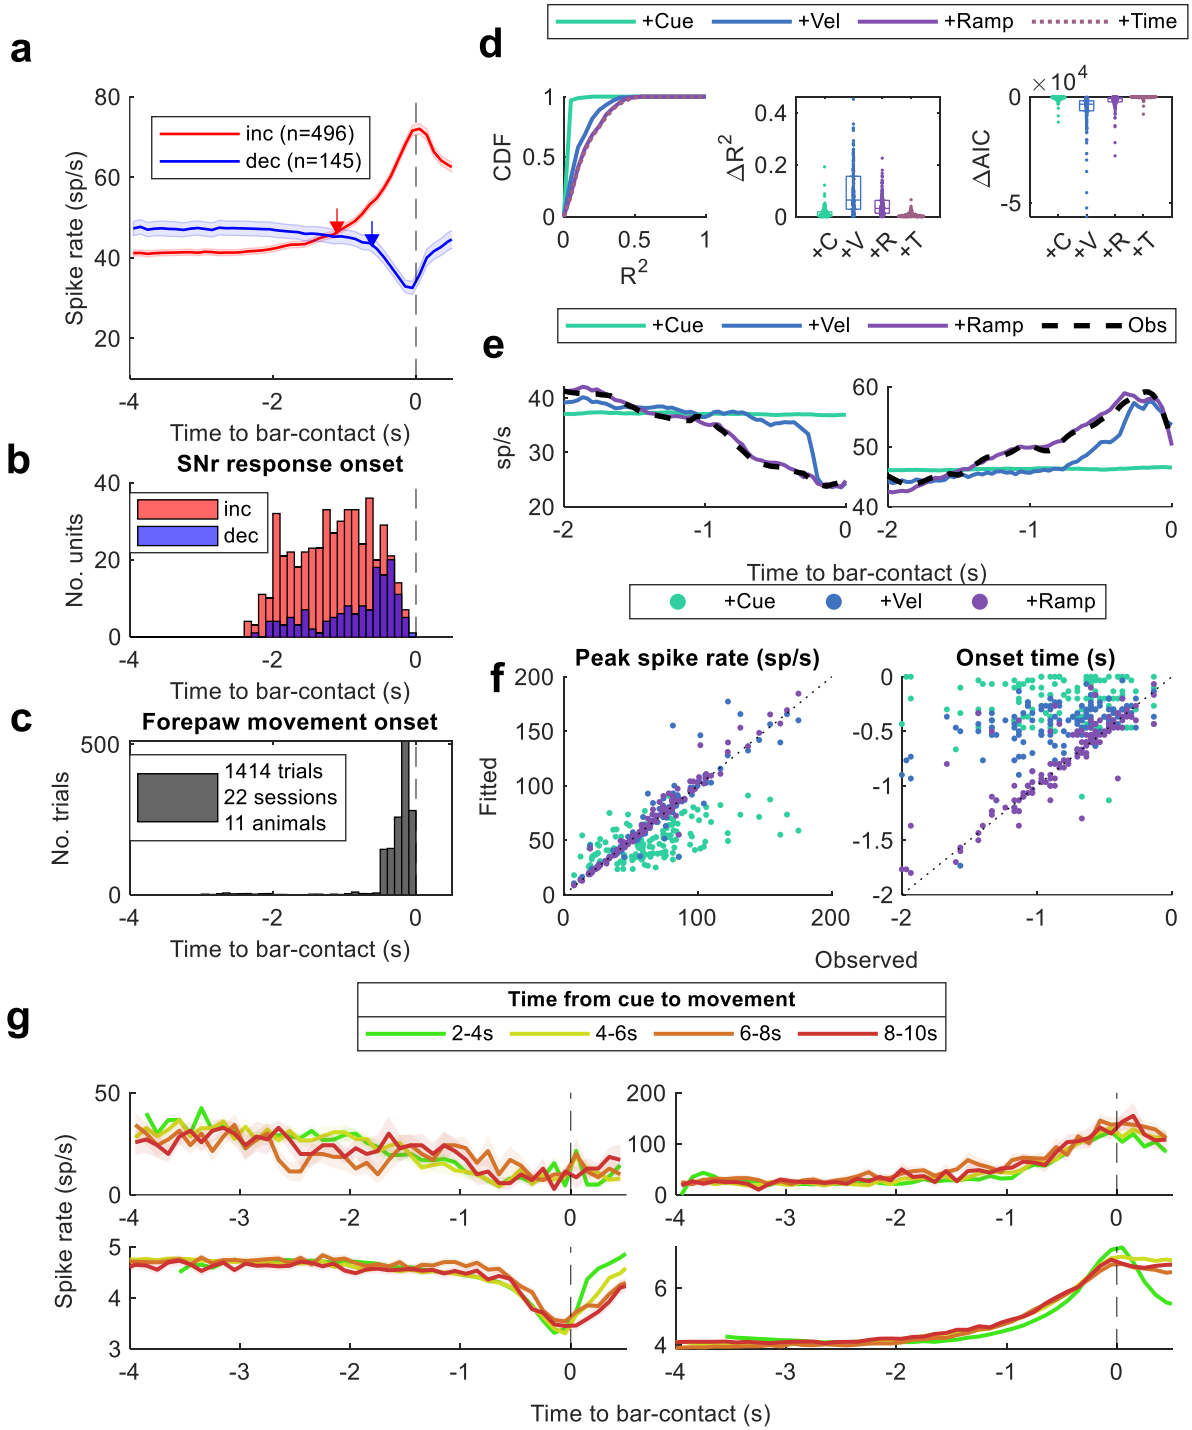

Figure 2. The timing of SNr neural activity in the self-timed reach task
a. Average PETH across increase units (red, n=496) or decrease units (blue, n=145); arrows indicate median neural response onset time (-1.10 s and -0.61 s for increase and decrease units, respectively). **b.** Distribution of neural response onset time for increase (red, n=496) and decrease (blue, n=145) units. **c.** Distribution of video-detected forepaw movement onset times relative to bar-contact. **d.** Improvement of model performance by iteratively adding predictors to the nested-GLM; **left:** cumulative distribution of explained variance (R^2); **middle:** additional variance explained (ΔR^2); **right:** change in Akaike information criterion (ΔAIC). Jittered dots represent individual SNr units (n=163), box plot represents the median & 25/75 percentiles. **e.** Fitted (colored traces) vs. observed (black dashes) movement-aligned PETH for two example SNr units. **f.** Fitted (y-axis) vs. observed (x-axis) PETH statistics; points represent 163 SNr units; colors represent different iterations of the nested-GLM; **left:** highest (or lowest) spike rate in the window [-0.5, 0] s before bar contact; **right:** onset latency (relative to bar-contact time) of neural response. **g.** PETH grouped by movement time; **top:** two example SNr units; **bottom:** grand average across 496 increase units (left) and 145 decrease units (right).

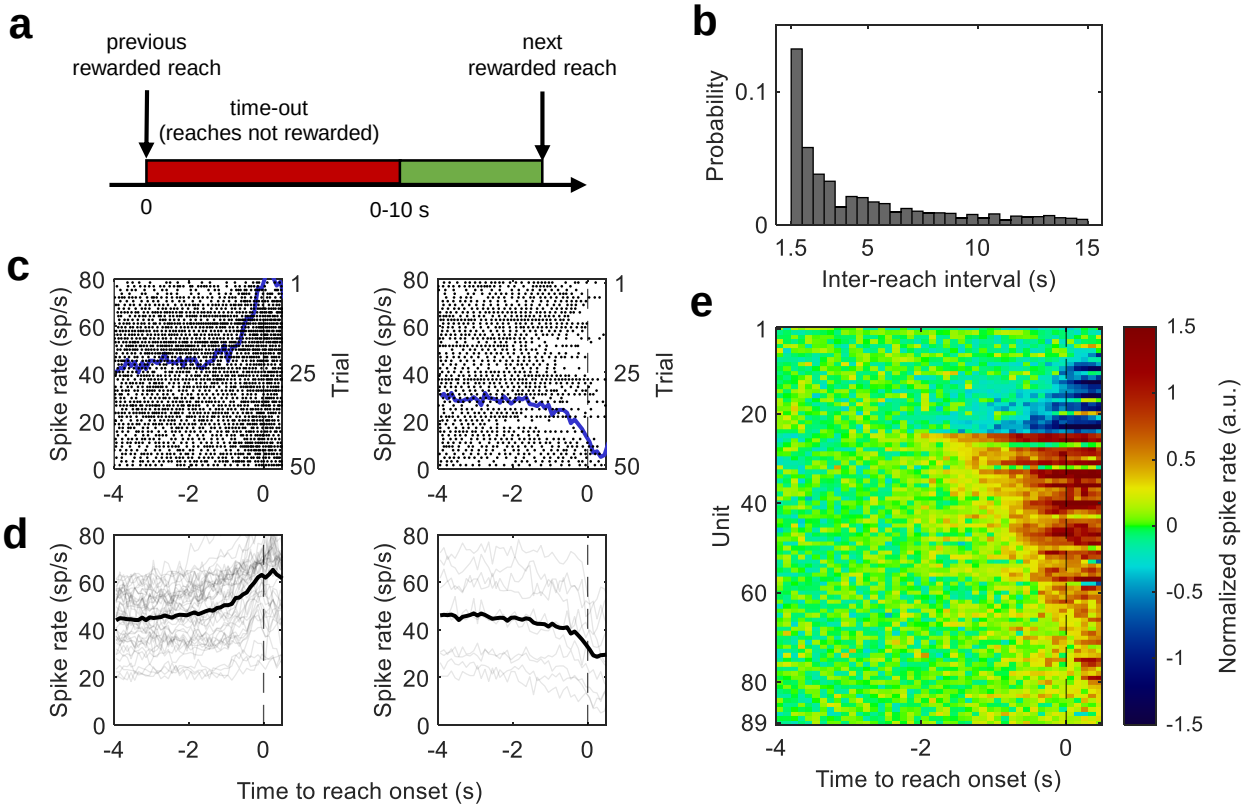

Figure 3. SNr neural activity in the spontaneous reach task
a. Time course of a spontaneous reach trial. The timeout is randomized between 0 and 10 s according to an exponential distribution. **b.** Distribution of inter-reach intervals (time between bar contacts; not shown beyond 15 s; n=1,164 trials from 3 animals, 8 sessions; individual sessions resemble the aggregate). **c.** PETH (blue trace) and spike raster (black dots) aligned to movement onset, for two example SNr units. For clarity, rasters show every 5th spike across 50 randomly selected trials. **d.** PETH of increase units (left, n=43) and decrease units (right, n=10). Dark black traces represent population average. **e.** Heat map visualization of PETHs for 89 units. Units are ordered by response sign and onset latency. **c-e:** trials with short (<6 s) inter-reach intervals are excluded.

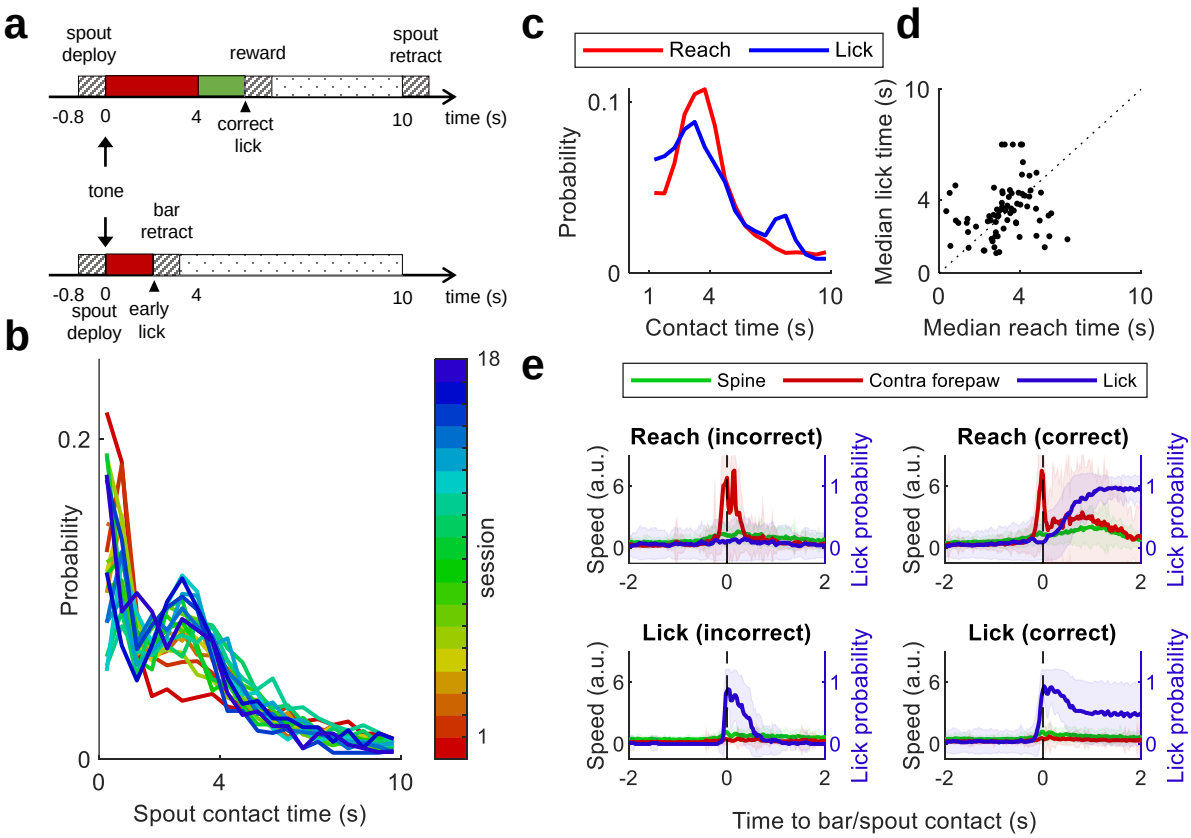

Figure 4. The self-timed lick vs. reach task

a. Example time course for a rewarded (top) or unrewarded (bottom) self-timed lick trial, contingent on lick time. For the reach task see Figure 1b. **b.** Distribution of movement times (i.e. first spout contact after cue), shown in different colors for the first 18 sessions that included self-timed lick trials (aggregate across 16 animals). **c.** Distribution of movement times across 11,238 self-timed reach trials (red) vs. 9,811 self-timed lick trials (blue). Reactive movements (made ≤1 s after cue) are not shown. **d.** Median reach time vs. median lick time relative to start cue for 79 sessions. **e.** Trial-averaged normalized movement speed of the spine (green) and the contralateral forepaw (red), as well as tongue protrusion probability (blue), derived from video analysis. Plots are shown separately for incorrect/unrewarded reach trials (n=71), correct/rewarded reach trials (n=185), incorrect/unrewarded lick trials (n=159), and correct/rewarded lick trials (n=176). Shaded area represent standard deviation.

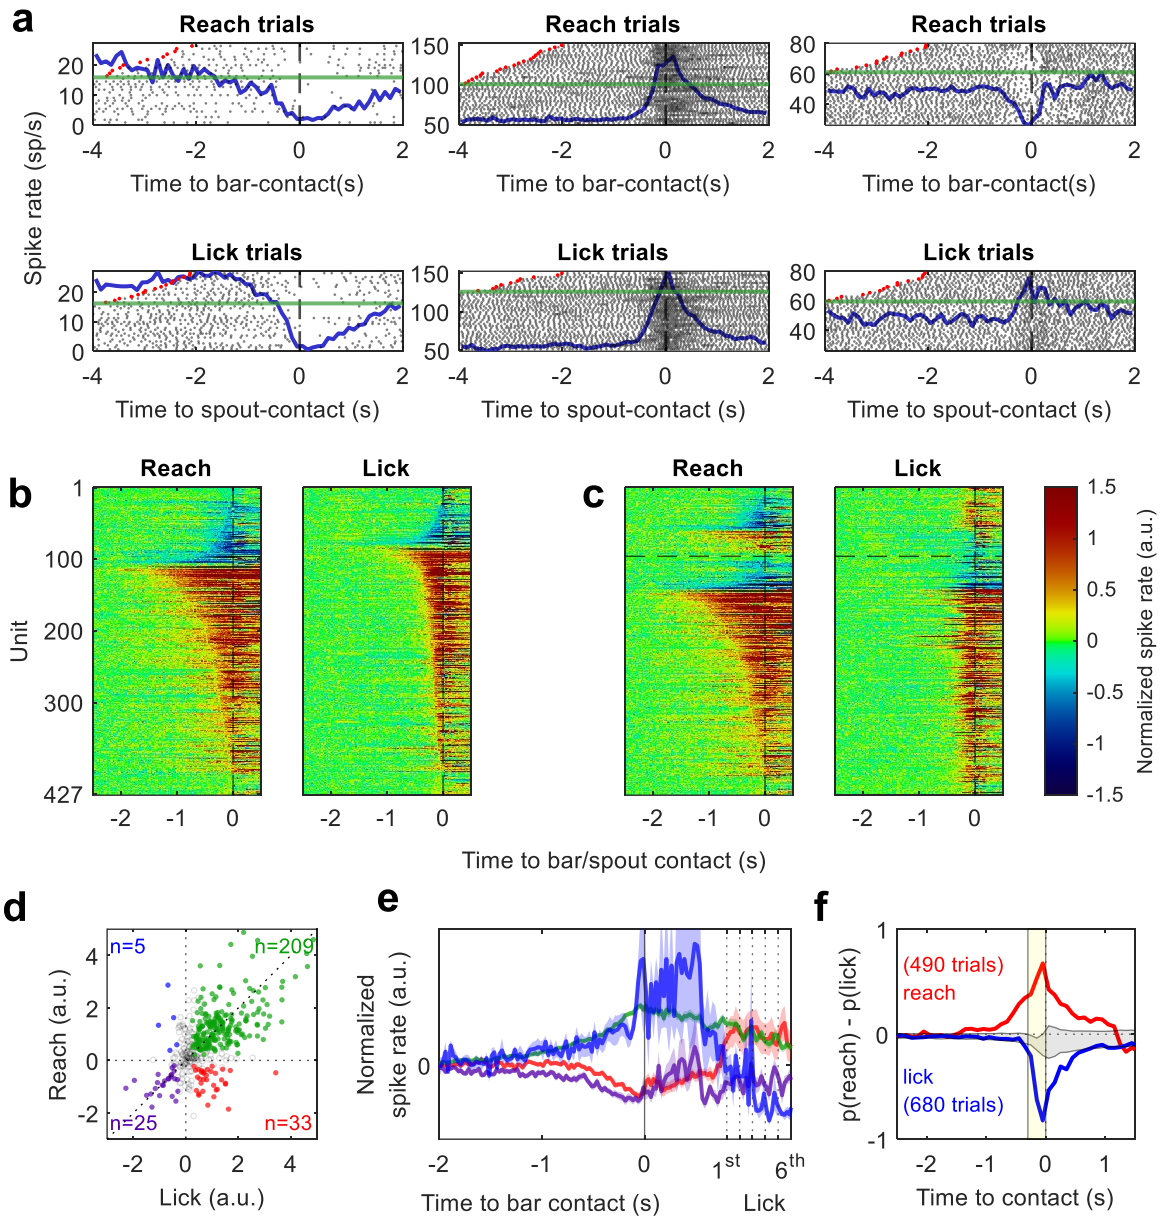

Figure 5. SNr neural activity in the self-timed lick vs. reach task.
a. Raster plots of spiking activity for three example SNr units during 40 self-timed reach trials (top) vs. 40 self-timed lick trials (bottom). Trials are sorted by movement time, with red dots indicating the start cue. The green line separates unrewarded trials (above) from rewarded trials (below). Black dots represent spikes (every 4th spike is shown for clarity). Blue traces represents PETHs. **b.** Heat map visualization of peri-reach (left) and peri-lick (right) PETHs for 427 SNr units. PETHs are normalized to baseline spike rates [-4, -2] s before movement. Units are ordered by sign and onset latency of the peri-movement response. **c.** Heat map visualization of peri-reach (left) vs. peri-lick (right) PETHs, with 427 SNr units sorted in the same order. The horizontal dashed line separates oppositely-modulated units (above) from congruent units (below). **d.** Normalized peri-movement activity ([-0.3s, 0] before bar/spout contact) during reach vs. lick: colored dots represent 272 units significantly modulated for both reach and lick, split into 4 subpopulations based on the sign of peri-reach vs. peri-lick responses; grey circles represent the remaining 155 units. **e.** Peri-movement neural activity during rewarded reach trials, averaged among the 4 subpopulations (as in **d**). Time after bar contact (0) is aligned to licking (i.e., 1st spout contact, etc.). **f.** LDA decoding of SNr population activity: decoded probability of reach subtracted by probability of lick is averaged among true reach trials (red) vs. true lick trials (blue). Yellow shaded area represents peri-movement training window: [-0.3, 0] s. Model is trained on SNr population activity from 490 reach trials and 680 lick trials (10 animals, 13 sessions, ≥10 SNr units per session, 293 total units). Grey shaded area represent 99% CI of null models trained on shuffled trials (10,000 permutations).

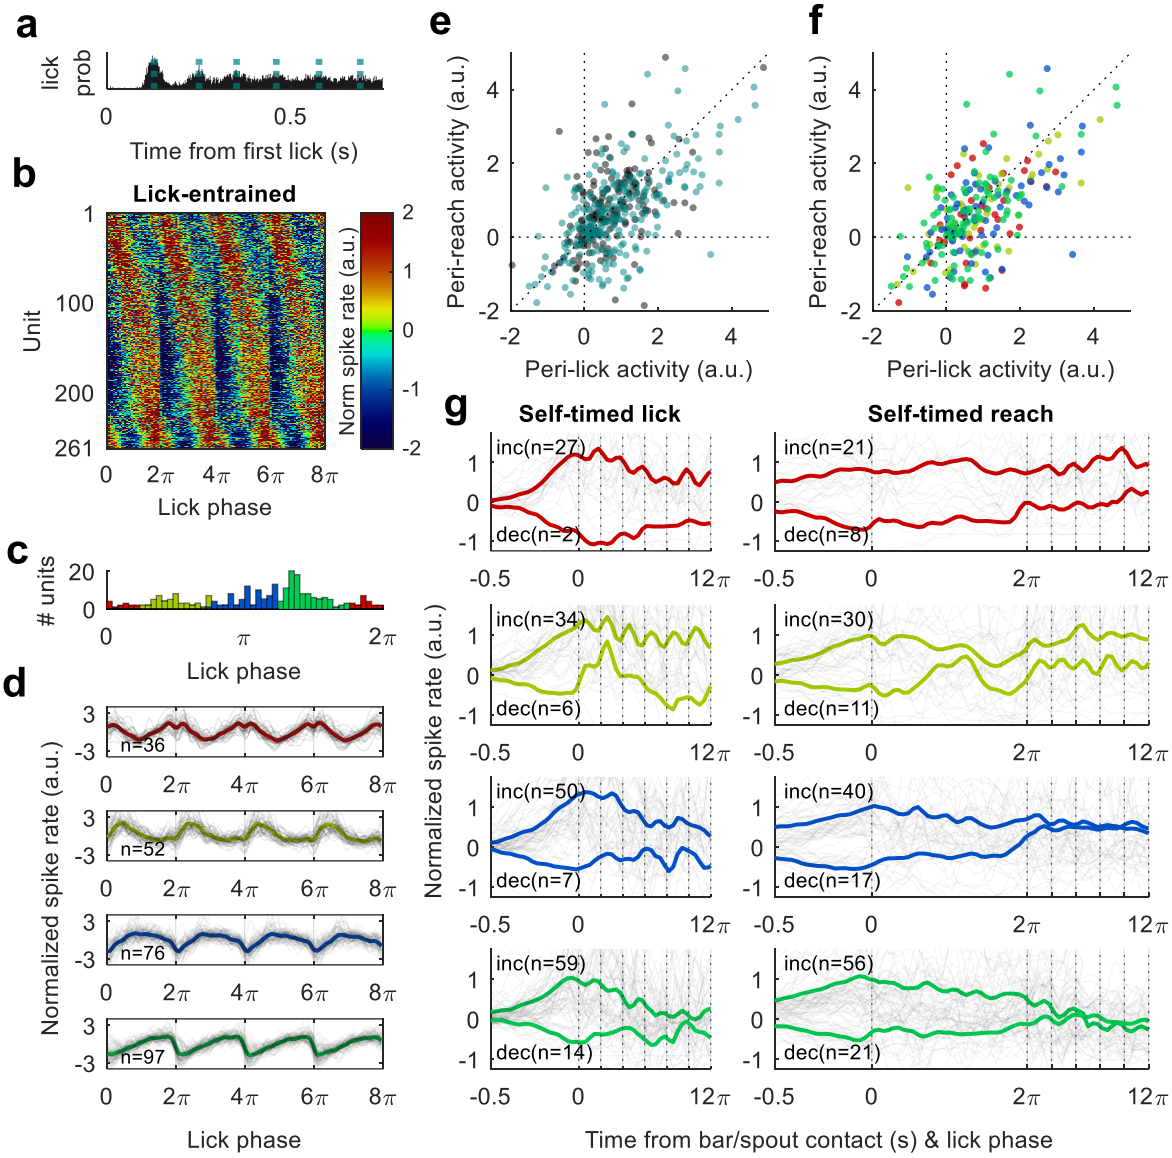

Figure 6. Lick-entrained SNr units

a. Lick histogram aligned to the first lick in each trial (n=33,337 trials); vertical dashed lines indicate estimated peaks in lick probability. **b.** Heat map visualization of normalized PETHs aligned to lick-phase (n=261 lick-entrained units). Units are ordered by preferred phase (0, 2*π*, etc. denote spout-contact). **c.** Distribution of phase preference amongst 261 lick-entrained SNr units. Anti-phasic and in-phase units are colored blue and red, respectively. **d.** Normalized spiking activity of lick-entrained SNr units, averaged across bouts of 4 lick cycles. **e.** Normalized peri-reach vs. peri-lick responses ([-0.3, 0] s before self-timed bar/spout contact) for 427 SNr units, of which 261 lick-entrained units are colored teal. **f.** Normalized peri-reach vs. peri-lick responses for 261 lick-entrained units, colored by phase preference (as in **c**). **g.** Peri-movement spike rates of lick-entrained and movement-modulated SNr units during rewarded self-timed lick (left) or rewarded self-timed reach (right) trials. Units are grouped by phase preference. Grey traces represent individual SNr units, colored traces indicate population averages (separate traces are shown for increase vs. decrease units). Spikes after the first lick are binned relative to the licking cycle. In rewarded self-timed reach trials, the first lick is delayed by ~0.8 s due to spout deployment.

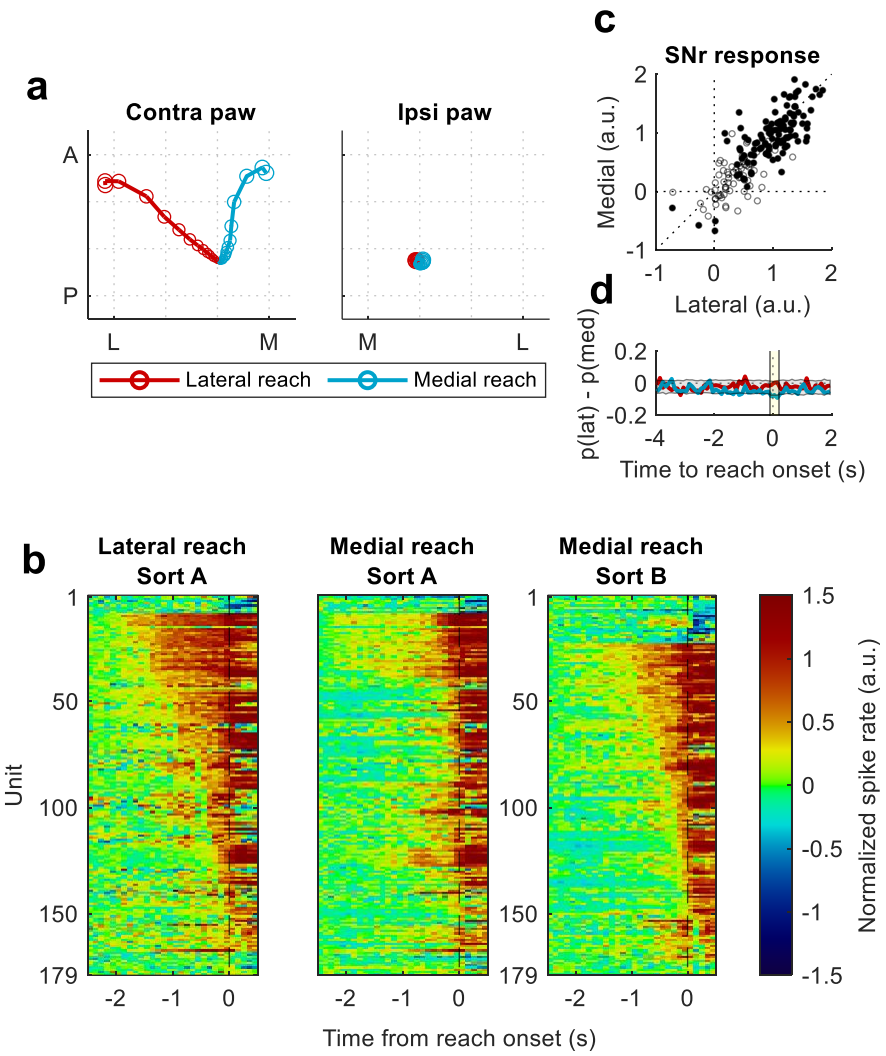

Figure 7. SNr neural activity in the directional reach task
a. Trial-averaged movement trajectories of the contralateral paw (left plot) and ipsilateral paw (right plot). Trajectories are centered relative to the resting position of the paw (i.e. 500ms prior to touch). Colors indicate whether the reach target was presented at the lateral (red, n=413 trials) or medial (blue, n=544 trials) position (3 animals, 10 sessions). Trials where animals reached with the ipsilateral paw or both paws (n=384 trials) were not included. **b.** Heat map visualization of normalized PETH for 179 SNr units during lateral reaches (left plot) vs. medial reaches (middle/right plots). Left two plots (sort A): units are ordered according to response sign/latency during lateral reach. Right plot (sort B): units are ordered according to response sign/latency during medial reach. **c.** Normalized spike rates around the time of reach onset (-100ms to 200ms) for 179 SNr units during lateral reaches vs. medial reaches, filled in dots (95/179) represent units with significant peri-movement responses for one or both reach directions (permutation test; $p<0.01$). **d.** LDA decoding of SNr population activity: decoded probability of lateral reach subtracted by probability of medial reach is shown for true lateral reach trials (red) vs. true medial reach trials (blue). Yellow shaded area [-0.1, 0.2] s represents training window. Model is trained on SNr population activity from 351 lateral reach trials vs. 424 medial reach trials (3 animals, 10 sessions, ≥10 SNr units per session, 165 total units). Grey shaded area represents 99% CI of null models trained on shuffled trials (10,000 permutations).

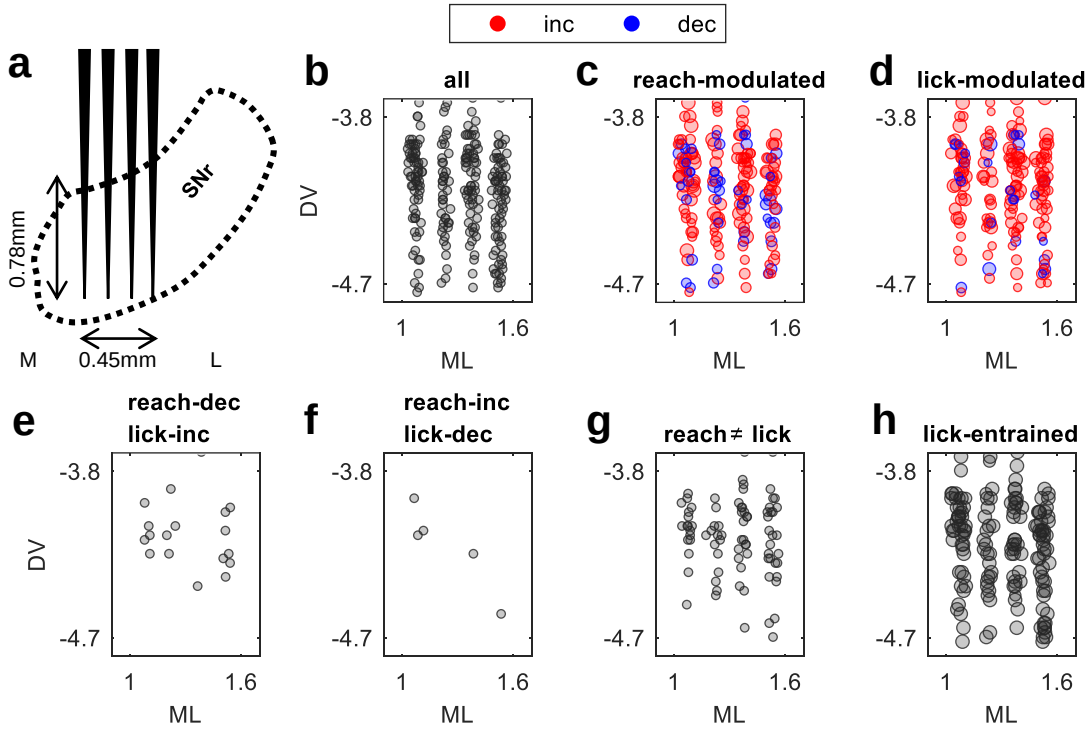

Figure 8. Location of reach vs. lick modulated SNr units

a. Illustration of SNr recording site with acute implantation of 4-shank 128-channel probes. **b.** Estimated location of SNr units recorded with 4-shank probes (n=242 units, 8 animals, 12 sessions). **c & d.** Location of SNr units significantly modulated during self-timed reach (n=193) and lick (n=186) tasks, respectively. Red and blue circles indicate increase or decrease units, respectively. Radius of circles represent amplitude of movement-related modulation. **e.** SNr units which decreased spike rates during reach but increased during lick (n=18), and **f.** vice versa (n=5). **g.** SNr units with significantly different spiking rates during reach vs. lick (n=81; $p<0.01$; permutation test). **h.** Lick-entrained SNr units (n=161).

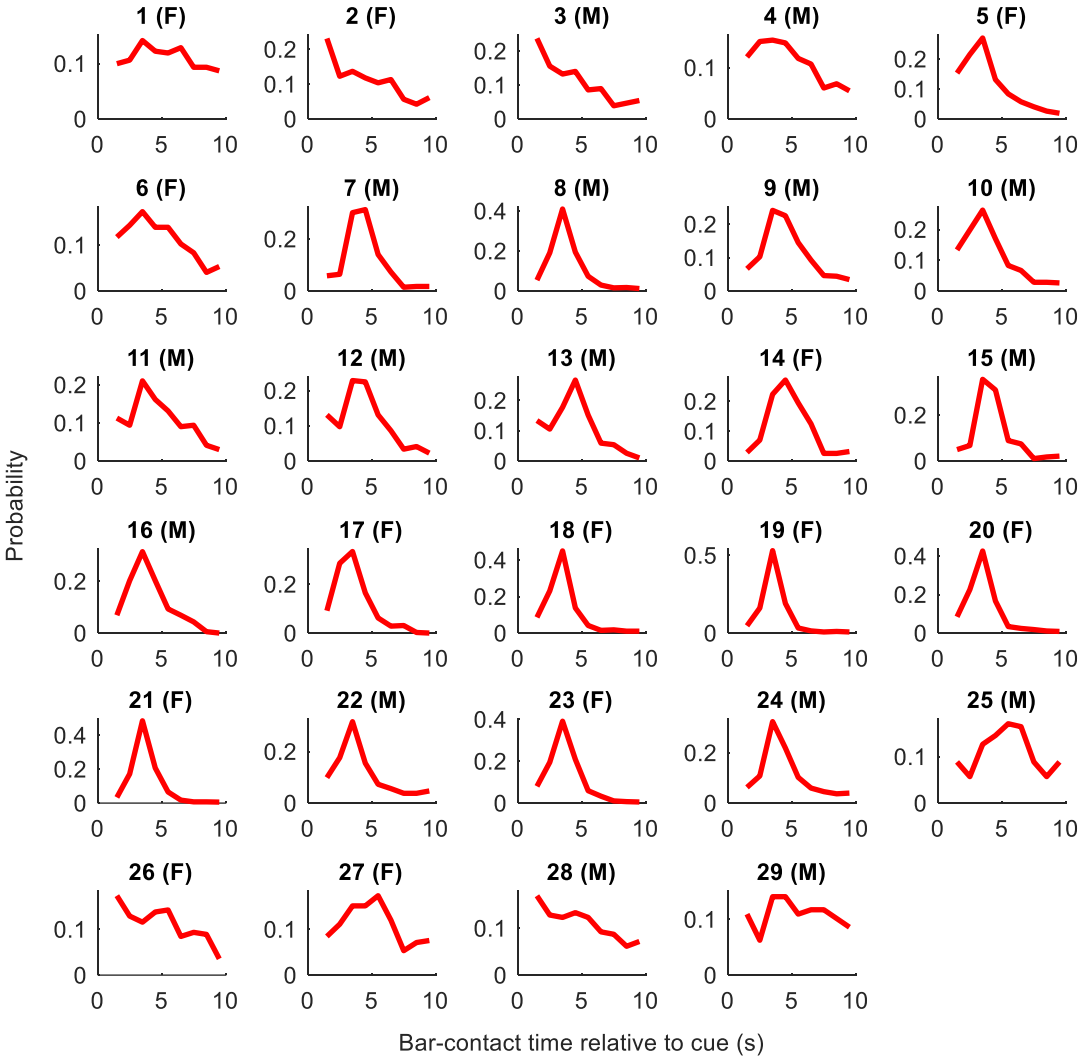

Figure S1. Movement time distribution for individual animals in the self-timed reach task

Movement-time histograms are generated for each of the 29 animals performing the self-timed reach task, aggregated across three best sessions after animals were fully trained. Reactive movements made before 1 s are not shown.

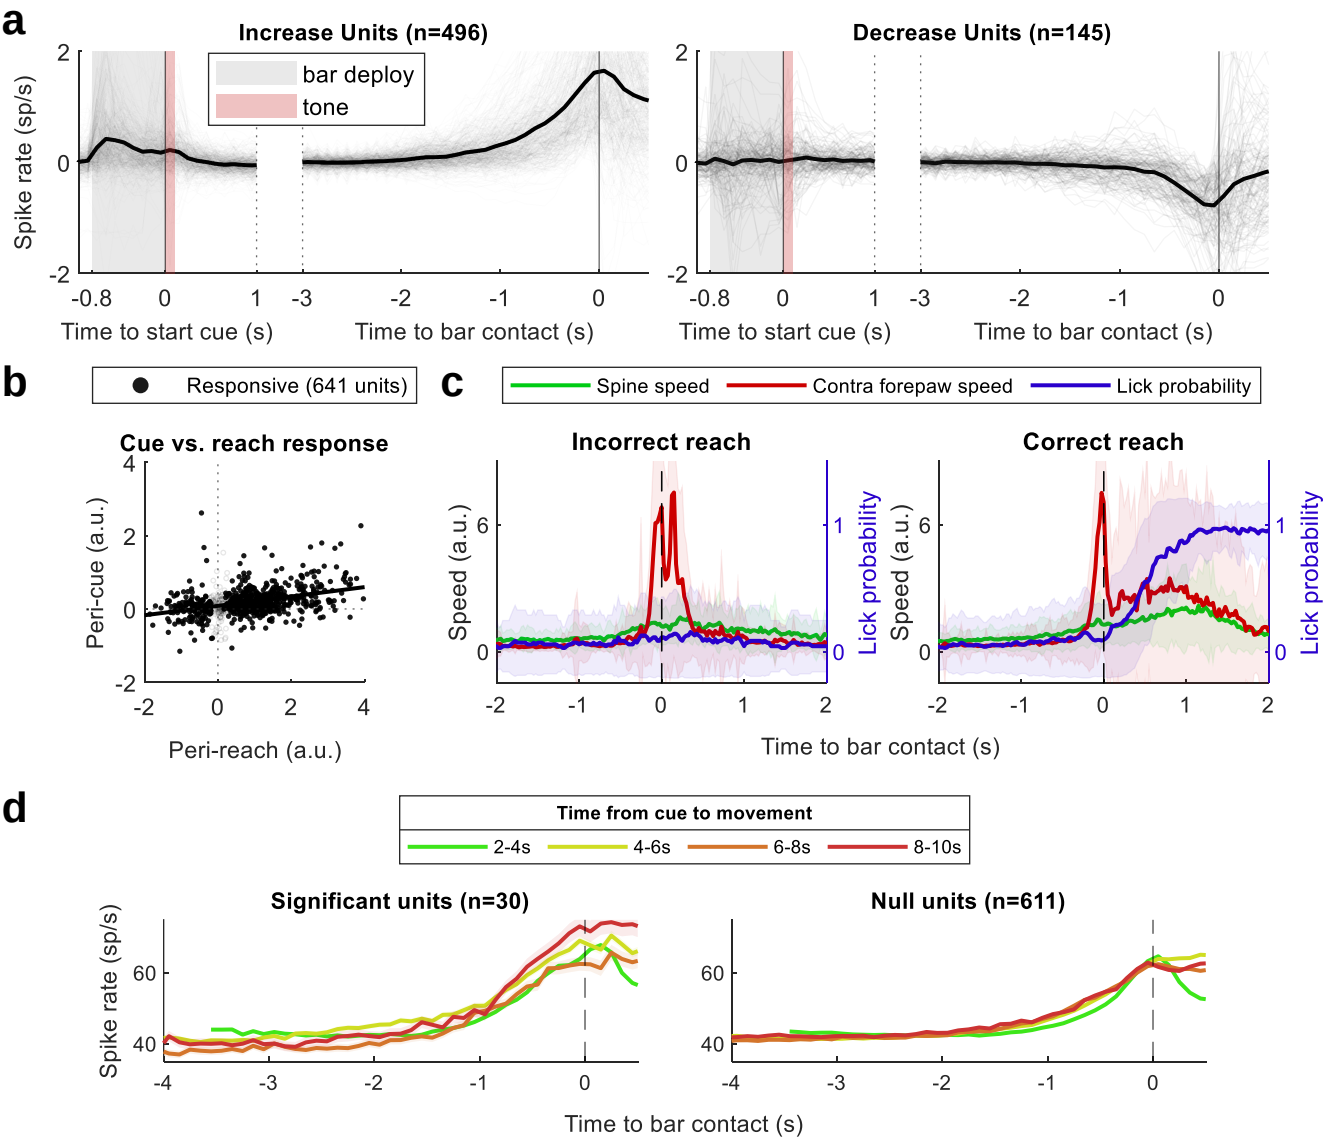

Figure S2.
a. Cue-aligned and movement-aligned PETH normalized to baseline (left: 496 increase units; right: 145 decrease units). Dark black trace represents population average. Grey and red shaded areas represent bar deployment (800 ms) and tone (100 ms), respectively. **b.** Normalized peri-cue activity ([-0.8, 0.1] s relative to tone onset) vs. peri-reach activity ([-0.3, 0] s relative to bar contact). Black dots represent movement-modulated units (n=641), grey circles represent non-modulated units (n=143). Black line represents linear regression (R^2=0.17). **c.** Trial-averaged normalized movement speed of the spine (green) and the contralateral forepaw (red), as well as tongue protrusion probability (blue), derived from video analysis. Plots are shown separately for incorrect/unrewarded (left, n=71) vs. correct/rewarded (right, n=185) reach trials. Shaded area represent standard deviation. **d.** Population average PETH during the self-timed reach task, grouped by movement time relative to the trial-start cue (as indicated by color). **Left**: population average across 30 SNr units whose pre-movement activity ([-2, 0] s before bar contact) differ significantly based on movement timing ($p<0.01$; permutation test); **Right**: population average across the remaining 611 SNr units whose pre-movement activity are similar regardless of movement timing.

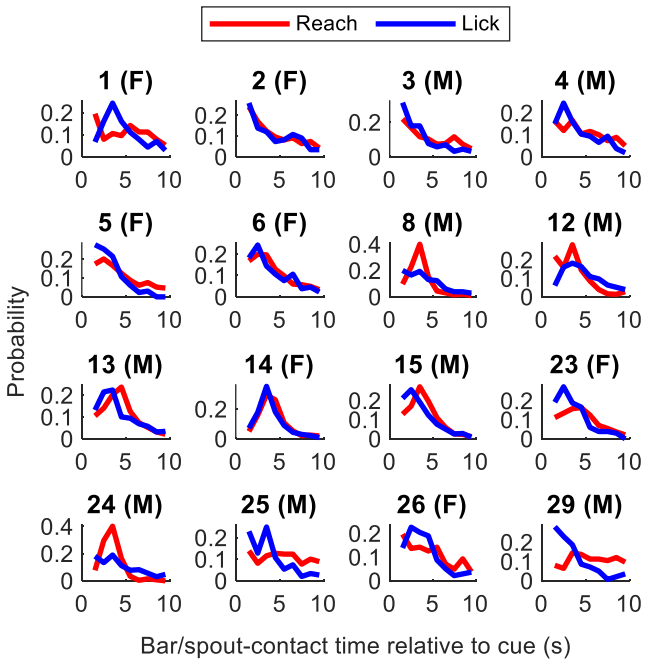

Figure S4. Movement time distribution for individual animals in the self-timed reach vs. lick task, for each of the 16 animals performing the self-timed reach vs. lick task, aggregated across three best sessions after animals were trained. Reactive movements made before 1 s are not shown. Reach times are shown in red; lick times are shown in blue.

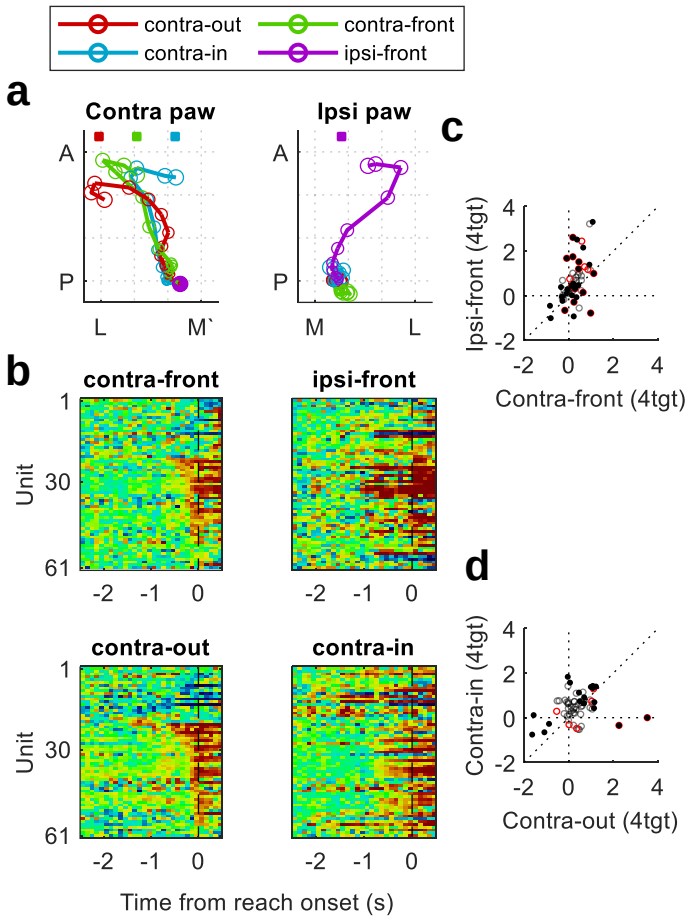

Figure S7. The four-direction forepaw reach task
We further trained two animals to perform the reach task with four possible target locations: lateral to, in front of, or medial to the contralateral paw, or in front of the ipsilateral paw (see Methods). We named these target locations “contra-out”, “contra-front”, “contra-in”, and “ipsi-front”, respectively. Animals preferred to use their contralateral forepaw for contralateral targets, and their ipsilateral forepaw for the ipsilateral target (rare trials in which animals used the opposite forepaw were excluded).
a. Trial-averaged movement trajectories of the contralateral forepaw (left) and ipsilateral forepaw (right) during the four-target reaching task. Trajectories are centered relative to the resting position of the paw (i.e. 500ms prior to touch). Colored squares indicate the latitude of the reach target. Trials where animals moved the wrong forepaw (i.e. not proximal to target) or both forepaws were excluded. **b.** Heatmaps showing normalized peri-movement activity of 61 SNr units during the four-target reach task. Units were sorted according to neural response latency during the “contra-front” target trials (for color scale see Fig. 8b). **c.** Normalized peri-reach activity ([-0.1, 0.2] s around reach onset) for 61 SNr units, during “contra-front” vs. “ipsi-front” target trials, where animals reached forward with the contralateral or ipsilateral paw, respectively. **d.** Normalized peri-reach activity ([-0.1, 0.2] s around reach onset) for 61 SNr units, during “contra-in” vs. “contra-out” target trials, where animals used the contralateral paw to reach inward or outward, respectively. In **c-d**, solid black dots indicate SNr units showing significant peri-movement responses in one or both movement types (permutation test; $p<0.01$; n=31/61, 19/61); red circles indicate SNr units whose peri-movement response amplitudes significantly differ across two movement types (permutation test; $p<0.01$; n=18/61, 7/61); grey circles represent the remaining unmodulated units.

Cohort	Self-timed reach	Self-timed reach vs. lick	2-target reach	4-target reach	Spontaneous reach
A (8M 3F)	291 units 113 sessions				
B (8M 10F)	493 units * 109 sessions	427 units † 76 sessions			
C (2M 3F)			179 units 10 sessions	61 units 3 sessions	
D (1M 2F)					89 units 8 sessions

Table S1. Number of SNr units and recording sessions, grouped by cohort (row) and experiment (column). † is a subset of *, as sessions with <20 self-timed lick trials were excluded.